



Sean Deuby on Enterprise Identity

# Windows IT Pro

A PENTON PUBLICATION

JUNE 2013 | WINDOWSITPRO.COM | WE'RE IN IT WITH YOU

## Windows 8 Security

Delete Active Directory  
User Accounts with  
Remove-ADUser

Get Started with Hyper-V  
in Windows Server 2012

Expand Tabs to  
Spaces in PowerShell

Understand  
System Center 2012  
Configuration Manager

# 1&1 DYNAMIC CLOUD SERVER

# PRICE CONTROL



AS LOW AS

**\$0-06**  
PER HOUR\*  
PLUS \$60 OFF  
FIRST MONTH



## COMPLETE COST CONTROL

- Full transparency with accurate hourly billing
- Parallels® Plesk Panel 11 included, with unlimited domains
- Sleep mode feature



## FULL ROOT ACCESS

- The complete functionality of a root server with dedicated resources



## MAXIMUM FLEXIBILITY

- Configure the Processor Cores, RAM and Hard Disk Space
- Add up to 99 virtual machines



## FAIL-SAFE SECURITY

- Redundant storage and mirrored processing units reliably protect your server



# 1&1



DOMAINS | E-MAIL | WEB HOSTING | eCOMMERCE | SERVERS

Call **1 (877) 461-2631** or buy online

**1and1.com**

\* Other terms and conditions may apply. Visit [www.1and1.com](http://www.1and1.com) for full promotional offer details. Customer is billed monthly for minimum configuration (\$0.06/ hour \* 720 hours = \$43.20/ month minimum). A \$60 credit (valid only for the first month) will be applied to your first month of service. Program and pricing specifications and availability subject to change without notice. 1&1 and the 1&1 logo are trademarks of 1&1 Internet, all other trademarks are the property of their respective owners. © 2013 1&1 Internet. All rights reserved.



# GRADUATE with **MORE**

Advancing your career in IT means you want more. At nonprofit WGU, we can offer you more—more real-world experience, more knowledge, and recognized IT certifications included as part of your curriculum.

WGU offers online bachelor's and master's degree programs in IT that are accredited, affordable, and respected by employers for their quality.

Visit [www.wgu.edu/itpro](http://www.wgu.edu/itpro) to learn how you can graduate with more with an IT degree from WGU.





# COVER STORY ▼

## Windows 8 Security 35

— Russell Smith

Systems administrators need to understand Windows 8's new security features to work effectively with the new operating system. This article discusses application containers, SmartScreen, IE Enhanced Protected Mode, and other low-level changes to Windows security.

### Features

#### 46 Getting Started with Hyper-V in Windows Server 2012

John Savill

#### 59 Expanding Tabs to Spaces in PowerShell

Bill Stewart

#### 63 Understanding System Center 2012 Configuration Manager

Orin Thomas

### Products

#### 74 New & Improved

### Interact

#### 69 Ask the Experts

### In Every Issue

#### 78 Ctrl+Alt+Del

#### 79 Advertiser Directory

#### 79 Directory of Services

#### 79 Vendor Directory

### Chat with Us



Facebook



Twitter



LinkedIn

# Columns



6 [Need to Know](#)

## **The Microsoft Transition, and How It Will Affect You**

Paul Thurrott



13 [Windows Power Tools](#)

## **Deleting Active Directory User Accounts with Remove-ADUser**

Mark Minasi



17 [Top 10](#)

## **Top 10 Essential Windows Server 2012 Keyboard Shortcuts**

Michael Otey



22 [Enterprise Identity](#)

## **Catching Up with the Cloud Security Alliance**

Sean Deuby



25 [What Would Microsoft Support Do?](#)

## **The Claims Rule Language in AD FS**

Joji Oshima

## Editorial

Editorial Director: Megan Keller  
Editor-in-Chief: Amy Eisenberg  
Senior Technical Director: Michael Otey  
Technical Director: Sean Deuby  
Senior Technical Analyst: Paul Thurrott  
IT Community Manager: Rod Trent  
Systems Management, Networking,  
Hardware: Jason Bovberg  
Scripting: Blair Greenwood  
SharePoint, Active Directory, Security,  
Virtualization: Caroline Marwitz  
SQL Server, Developer Content:  
Megan Keller  
Managing Editor: Lavon Peters  
Editorial SEO Specialist: Jayleen Heft

## Senior Contributing Editors

David Chernicoff, Mark Minasi,  
Tony Redmond, Paul Robichaux,  
Mark Russinovich, John Savill

## Contributing Editors

Alex K. Angelopoulos, Michael Dragone,  
Jeff Felling, Brett Hill, Dan Holme,  
Darren Mar-Elia, Eric B. Rux,  
William Sheldon, Curt Spanburgh,  
Bill Stewart, Orin Thomas, Douglas Toombs,  
Ethan Wilansky

## Art & Production

Senior Graphic Designer: Matt Wiebe  
Director of Production: Dylan Goodwin  
Group Production Manager:  
Julie Jantzer-Ward  
Project Manager: Adriane Wineinger  
Graphic Specialists: Karly Prickett &  
Ashley Lawson

## Advertising Sales

Technology Market Leader: Peg Miller  
Key Account Director:  
Chrissy Ferraro • 970-203-2883  
Account Executives:  
Megan Key • 970-203-2844  
Barbara Ritter • 858-367-8058  
Cass Schulz • 858-357-7649

## Client Services

Senior Client Services Manager:  
Michelle Andrews • 970-613-4964  
Ad Production Coordinator: Kara Walby

## Marketing & Circulation

Customer Service • 800-793-5697  
Senior Director, Marketing Analytics:  
Tricia Syed

## Technology Division & Penton Marketing Services

Senior Vice President: Sanjay Mutha

## Corporate

Chief Executive Officer:  
David Kieselstein  
Chief Financial Officer/Executive Vice  
President: Nicola Allais



## List Rentals

MeritDirect  
333 Westchester Avenue,  
White Plains, NY 10604

## Reprints

Reprint Sales:  
Wright's Media • 877-652-5295

*Windows IT Pro*, June 2013, Issue No. 226,  
ISSN 1552-3136. *Windows IT Pro* is published monthly by  
Penton Media, Inc. Copyright ©2013 Penton Media, Inc.  
All rights reserved. No part of this publication may be  
reproduced or distributed in any way without the written  
consent of Penton Media, Inc.

*Windows IT Pro*, 748 Whalers Way, Fort Collins, CO 80525,  
800-621-1544 or 970-663-4700. Customer Service:  
800-793-5697.

We welcome your comments and suggestions about the  
content of *Windows IT Pro*. We reserve the right to edit all  
submissions. Letters should include your name and  
address. Please direct all letters to [letters@windowsitpro.com](mailto:letters@windowsitpro.com). IT pros interested in writing for *Windows IT Pro* can  
submit articles to [articles@windowsitpro.com](mailto:articles@windowsitpro.com).

Program Code: Unless otherwise noted, all programming  
code in this issue is ©2013, Penton Media, Inc., all rights  
reserved. These programs may not be reproduced or  
distributed in any form without permission in writing from  
the publisher. It is the reader's responsibility to ensure  
procedures and techniques used from this publication are  
accurate and appropriate for the user's installation. No  
warranty is implied or expressed.

Windows®, Windows Vista®, and Windows Server®  
are trademarks or registered trademarks of Microsoft  
Corporation in the United States and/or other countries  
and are used by Penton Media, Inc., under license from  
owner. *Windows IT Pro* is an independent publication  
not affiliated with Microsoft Corporation. Microsoft  
Corporation is not responsible in any way for the editorial  
policy or other contents of the publication.

# Windows IT Pro

# The Microsoft Transition and How It Will Affect You



## Paul Thurrott

is senior technical analyst for *Windows IT Pro*. He writes the SuperSite for Windows, a weekly editorial for *Windows IT Pro UPDATE*, and a daily Windows news and information newsletter called *WinInfo Daily UPDATE*.

Email



Twitter



Website



**T**his month, I'd like to evaluate where Microsoft is in its transition to a maker of devices and services. This is the biggest transition in the company's history, one that will affect customers and users. Using Microsoft's most recent earnings release, let's rate the company's progress and determine which strategic transitions might affect you.

Microsoft announced its results for the first calendar quarter of 2013—what it calls its third quarter of fiscal year 2013—in April 2013. The results were stunning and unexpected: Microsoft reported net income of \$6.06 billion on record revenues of \$20.4 billion. That the company did this during the worst downturn in PC history is perhaps a testament to both the diversity of its business and the success of the ongoing transition of its core businesses.

Much of the information Microsoft provided in the press release, earnings conference call, and other related material focused on what I call “momentum marketing.” Normally, this is information I'm not particularly keen on, and it's been my experience that when Microsoft talks momentum this means it has nothing interesting to say. But in the case of quarterly financial announcements, in particular this one, product momentum is actually important. And that's because Microsoft is changing.

As you know, there are two aspects to the new Microsoft: devices and services. “Devices” refers to the devices Microsoft makes—a small lineup, admittedly, that includes the Surface tablet as well as the Xbox console—but also, indirectly, to a growing and diverse range of heterogeneous devices (based on Android and Apple iOS, for the most part) that Microsoft supports with mobile apps.

“Services” also has two meanings: online services such as Bing that are standalone services with no corresponding on-premises

alternative, and those services that represent what I see as Microsoft's future: its traditional on-premises client and server products remade as online services. In the case of server products, the transition is easy enough to understand: Exchange, SharePoint, and Lync can (conceptually at least) be hosted in the cloud as easily as they can in your own data center, or that of a partner. In the case of clients, Microsoft still needs to deliver bits of code to PCs and devices.

But what's changed—in products like Windows and Office, for example—is how these bits are deployed and, going forward, how they're serviced. This happens, as it does with online services, with smaller and more frequent updates than in the past.

I wrote about this transition last month in “[Windows 8 Blue and How Microsoft Is Delivering Software as a Service](#),” describing it in general and examining how [Windows 8](#), in particular, will be updated going forward. Now, I'd like to evaluate the broader Microsoft product families called out in Microsoft earnings, see where the firm is in transitioning those products to match its new focus, and discuss where you should focus your own platform transitions.

## Office

Microsoft's biggest business is Office, not Windows. In Q1, the Microsoft Business Division launched what the firm calls “the new Office,” or Office 2013, the first version of the suite to utilize *Click to Run*, meaning that the software can now be delivered and serviced like an online service. This new Office looks and works a lot like its predecessor, so training costs will be nonexistent to low; on disk, the product is virtually identical. It's possible to deploy the new Office the old way (legacy MSI-based installers are still available) but there will be a lot to learn for those who want to move forward.

This shouldn't quell enthusiasm. While discussion around Microsoft is focused on Windows 8 these days, the new Office is as big a leap forward as is Windows 8. But in this case, the important changes are all on the infrastructure side. Microsoft's historic “better together”



vision is realized when organizations roll out the new Office servers—Exchange, [SharePoint](#), and Lync 2013—in either on-premises or hosted forms alongside Office 2013. For users, the integration between these products—including such simplicities as tying SharePoint-based document libraries to the applications as default save and open locations—offers huge benefits.

And hosted Office is here to stay. One in four of Microsoft's enterprise customers now uses [Office 365](#) in some capacity and the business is on a \$1 billion annual revenue run rate.

With this new generation of Office, the hosted versions of the products will be updated frequently with new features and functionality. This transition of Office to a service is profound, and the division's plans to update the products—quarterly, as with Office 365—are both transparent and sound. Now is the time to begin thinking about when and if your business's on-premises Office servers can be moved to cloud services.

## Server

Microsoft's Server and Tools division is responsible not only for Windows Server but also SQL Server, System Center, and the ever-growing Azure business. As with Office, the transition to the cloud is now in full swing, and this past quarter Microsoft delivered Windows Azure Active Directory alongside the new version of Office 365. As Microsoft puts it, the firm is now “the only cloud provider that can offer customers a comprehensive hybrid cloud solution that integrates existing IT infrastructure with all the benefits of the public cloud.”

That said, infrastructure doesn't transition quickly or lightly. Over the past decade, we've seen ever-increasing Microsoft support life cycles on both the client and server, and a confluence of distrust in the cloud—especially the public cloud—and a fairly regular and impressive increase in the quality and functionality of Microsoft's Windows Server product line, in particular. With Windows Azure now beginning to offer a full spectrum of services, we're entering a

period of early evaluation only. Microsoft added 50 new services to Azure over the past year alone, and Windows Azure Infrastructure Services became generally available in April. Most businesses would benefit more from moving email and communications infrastructure to the cloud (Office 365) before moving PC and device management (System Center/Intune) or identity, storage, and related infrastructure (Windows Server/Azure).

## Windows

Windows 8 is evolving into the most controversial Windows release of all time, as much the result of our highly connected echo-chamber times as any functional weirdness. That Microsoft declined to release Windows 8 license sales figures with its quarterly financial results, is, I think, telling. From 2009 to date, the firm has consistently sold about 20 million Windows licenses a month, and the suspicion is that Q1 2013 was the first quarter in which that didn't happen. One can only conclude that Windows 8 confuses too many in a time in which viable mobile computing options—various tablets and smartphones—exist. And no business would ever broadly roll out the “touch first” Windows 8 this early in its life cycle.

The Windows division isn't offering the same clarity as Office when it comes to updates. Microsoft CFO Peter Klein did note that there is “a new, accelerated pace for updates and innovations,” which includes the ongoing release of updates to the built-in “Metro-style” mobile apps that ship with Windows 8 as well as, interestingly, a routine monthly updating of the firmware in its own Surface tablets (of which there are Windows 8 and Windows RT variants).

Klein also noted that “the next version of Windows, code-named ‘Windows Blue,’ [includes] further advances to the vision of Windows 8 as well as responds to customer feedback.” That response to customer feedback—some of which surely included customers simply ignoring Windows 8 for the time being due to its many defects—is important, I think.

But it doesn't change my guidance: Except when Windows 8 answers a real need—including vertical deployment scenarios in which a tablet device actually makes sense—no business of any size should actively be evaluating a Windows 8 deployment at this time. Microsoft will reveal more information about the Blue release in June, and you can determine whether the change it's making will overcome the training costs associated with Windows 8. My guess is that they will not.

Also pushing out the timetable on Windows 8 is the availability of new Intel chipsets that will help this fledgling OS make more sense on the mobile devices it was designed for. Microsoft notes that Intel's fourth-generation Core processor, code-named Haswell, will help enable new devices that combine performance benefits with power savings, and will be shipping in volume by mid-2013. And later this year, a new generation of devices based on Intel's Bay Trail Atom processors will drive prices down for low-end touch devices. Today, the cheapest touch devices are in the \$500 range. Intel would like to see them sell for as low as \$200.

Windows will transition to “a new era of mobile computing,” as Microsoft puts it, but that transition will take time, and Windows 8 will get better as we move forward. You should ride this one out.

## Online Services

Despite its name, the Online Services division won't affect most businesses directly. This is Microsoft's Bing business and online advertising efforts, not its online services such as Office 365 or Azure.

## Xbox and Windows Phone

Microsoft's Entertainment and Devices division is responsible for a curiously diverse range of products, which can be simplified down to “Xbox and then everything else.” Microsoft's struggling Windows Phone platform falls under this division. However, Microsoft has nothing but vague pronouncements on this front.

“Momentum with Windows Phone continues to build,” Microsoft general manager Chris Suh said. “The devices, now available at a broad range of price points, are receiving great reviews, and carrier support continues to grow. We now have over 10 percent share in several countries, but realize that there is still a lot of work ahead to break through in some key markets. With growing awareness of Windows Phone, and sustained innovation from our hardware partners, we feel well positioned to continue our momentum.”

The double appearance of the word momentum in that quote should trouble you: Windows Phone has roughly 4 percent market share worldwide and is utterly crushed by Android (particularly Samsung) and Apple iPhone in all of the markets that matter.

That said, those who have made big bets on Microsoft platforms across the board would do well to consider Windows Phone. No other smartphone platform offers the same level of integration with Microsoft’s enterprise-oriented technologies, including a fairly full-powered mobile Office suite (with an especially good version of OneNote), full support for Exchange ActiveSync policies, support for BitLocker, Information Rights Management (IRM), and more.

This is one area where I feel businesses aren’t moving quickly enough. If you’re using the Microsoft stack at all, Windows Phone needs to be on your radar. ■





# PASS SUMMIT 2013

October 15-18, 2013 | Charlotte, NC  
Charlotte Convention Center

## This is Community.

PASS Summit is the world's largest and most intensive technical training conference for Microsoft SQL Server and BI professionals. But more than that, it's your conference - planned and presented by the SQL Server community for the SQL Server community.

Join us in Charlotte and get the top-notch training, technical tips and tricks, and networking and connections you need to take your SQL Server skills to the next level.

REGISTER NOW!

[www.sqlpass.org/summit](http://www.sqlpass.org/summit)



Best  
conference  
I have  
attended.



The sessions  
were outstanding  
and well worth  
my time.



# Deleting Active Directory User Accounts with Remove-ADUser

Account deletion is a task to approach very carefully, and PowerShell gives us a safe method

In “Use Get-ADUser to Determine Who Has Never Logged On,” I showed you how to use PowerShell to find all Active Directory (AD) users who meet some criterion—they’re locked out, they haven’t logged on in a certain time period, they have first names starting with “J,” and so on. In “Doubling Up Active Directory PowerShell Cmdlets,” I showed you the commands that let you do something to those folks, such as unlock their account, disable or enable the account, change an account attribute, and so on. What I haven’t tackled yet is account deletion.

If you’ve read even a few of my previous columns about AD’s PowerShell cmdlets, you’ve seen that PowerShell cmdlets glue a “verb-ish” word to a particular noun to create commands such as *get-aduser*, and you know that the relevant noun for AD users is *ADuser*. You’ve probably also seen that PowerShell tries to restrict itself to a fairly short list of verb-ish words, and that the big four are *new* (which creates PowerShell objects), *get* (which displays PowerShell objects that meet some set of criteria), *set* (which lets you modify some aspect of an existing PowerShell object), and *remove* (which is PowerShell’s verb for *delete*). Knowing all that, you’ve probably already guessed that the command to delete a user account is *remove-aduser*. The command is quite simple in its most basic form:



## Mark Minasi

is a senior contributing editor for *Windows IT Pro*, an MCSE, and the author of 30 books, including *Mastering Windows Server 2008 R2* (Sybex). He writes and speaks around the world about Windows networking.



Email



Twitter



Website

```
remove-aduser identity
```

In that cmdlet, *identity* works as you’ve already seen it work in *get-aduser* and *set-aduser*: It will take a DN (cn = AprilJones,CN = users,dc = bigfirm,dc = com), a SID (S-1-5-21-941799636-306785290-3997453140-1106), an object GUID (c5959b71-61f9-4497-81a6-c147639a33b0), or a SAM account name (AprilJones). *Remove-aduser* is different from most AD cmdlets, however, in that it requires a confirmation. Try deleting someone, and you’ll see something like this:

```
Are you sure you want to perform this action?
Performing operation "Remove" on Target "CN=AprilJones,CN=Users,
    DC=bigfirm,DC=com".
[Y] Yes [A] Yes to All [N] No [L] No to All [S] Suspend [?] Help
    (default is "Y"):
```

Press Enter, and the deed is done. But that’s not particularly helpful, because many admins wouldn’t use *remove-aduser* to delete just one user account; instead, you might use it in a one-liner, as in

```
Search-ADAccount -AccountDisabled -UsersOnly | remove-aduser
```

That command would find all the disabled users and delete their accounts. (Please don’t run this command on your production network!) And if you plan to type that line on a domain with 1,000 disabled accounts, get ready to wear out your Enter key!

To tell *remove-aduser* not to require a confirmation, you could—well, what could you do? PowerShell is somewhat bipolar about this. PowerShell cmdlets tend to have some sort of *Are you sure?* prompt built into the “dangerous” commands, and that’s a good thing. What’s odd, however, is that to inform PowerShell that you know what you’re doing and that that you don’t need to press Enter, you sometimes use *-force*, as in *stop-process*:

```
stop-process -name "SomeService"
```

If “SomeService” is a process running under an account that isn’t you, as with most services, PowerShell will ask for confirmation, as *remove-aduser* did. As with *remove-aduser*, you can tell PowerShell not to ask for that confirmation, but only with *-force*:

```
stop-process -name "SomeService" -force
```

Try adding *-force* to the end of that *remove-aduser* command, however, and you’ll get the error message *A parameter cannot be found that matches parameter name ‘force’*—which, I have to say, is somewhat clearer than many PowerShell error messages. Instead, you can type the following to perform a deletion without any remonstrations:

```
remove-aduser AprilJones -confirm:$false
```

Personally, I like this second approach better because it identifies *confirm* as an internal flag that, when set to *\$true*, tells PowerShell to make a final check with you before doing something, and, when set to *\$false*, tells PowerShell to remain mum and just do what you told it to do. So always remember: If *-confirm:\$false* doesn’t stop PowerShell confirmations, *-force* will, and vice versa. (And like most of the “dangerous” commands, *remove-aduser* has the *-whatif* parameter, which reports on what it would have done without *-whatif*.)

How can you be sure that April is gone? Well, a simple

```
get-aduser AprilJones
```

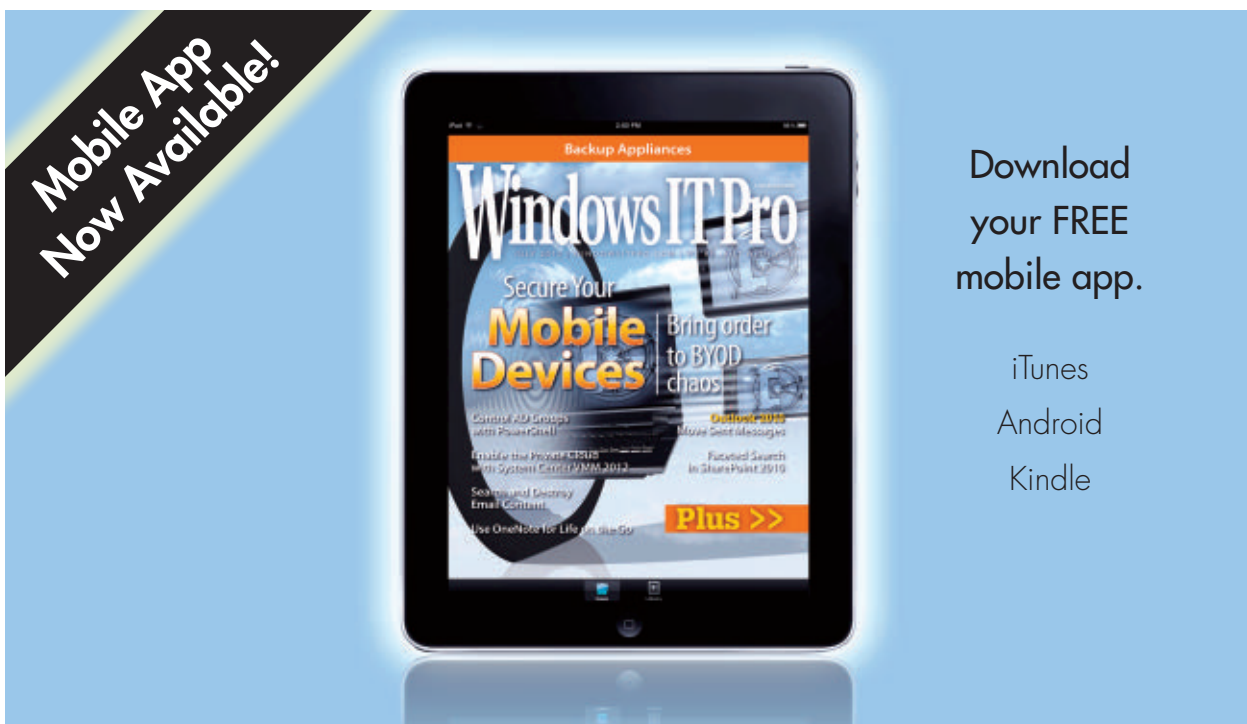
would do the trick, or you could (if you’ll pardon the gruesome expression) examine the corpse. As you probably know, deleting a user account in AD doesn’t actually erase the user object from AD but instead clears most of the object’s attributes and marks it as deleted,



creating a tombstone object that AD keeps around for a number of days (180, by default). To see tombstones, you need a cmdlet that's a bit more powerful than *get-aduser*. You need *get-adaccount*. Its syntax is like *get-aduser*'s, but it has an extra parameter, *-includedeletedobjects* (which can, fortunately, be shortened to *-inc*), that shows the otherwise-hidden tombstones. Search for April's remains with this:

```
get-ADObject -inc -f {samaccountname -eq "AprilJones"}
```

Run that query, and you'll see that April isn't quite gone yet, and in fact you might be able to undelete her account to some degree. But that's next month's topic! ■



**Mobile App  
Now Available!**

Download  
your **FREE**  
mobile app.

iTunes  
Android  
Kindle

# Top 10 Essential Windows Server 2012 Keyboard Shortcuts

See how keyboard shortcuts can boost your productivity with Microsoft's newest server OS

**M**icrosoft's latest server OS is more dependent on the keyboard than any version of Windows Server that I can ever remember using—and I've used them all. If you're working with [Windows Server 2012](#), here are 10 keyboard shortcuts that you definitely need to know about. And if you manage Windows Server systems remotely through Remote Desktop, remember to set the *Apply Windows key combinations* RDP setting to *On the remote computer*.

## ① Windows key+X

The mother of all Windows Server 2012 keyboard shortcuts, the Win + X combination is an absolute necessity. It lets you access:

- Programs and Features
- Power Options
- Event Viewer
- System
- Device Manager
- Disk Management
- Computer Management
- Command Prompt
- Task Manager
- Control Panel



## Michael Otey

is senior technical director for *Windows IT Pro* and *SQL Server Pro* and author of *Microsoft SQL Server 2008 High Availability with Clustering & Database Mirroring* (McGraw-Hill).



Email

- File Explorer
- Search
- Run

It seems like Microsoft put everything it couldn't find a better place for under the Win + X shortcut.

## ② Windows key+I

Probably the next most important keyboard shortcut is the Win + I combination. The Win + I keyboard shortcut opens the new Windows Server 2012 Settings menu where you can access Control Panel, Personalization, Server Info, and Help options. It also provides access to Network options, Audio options, Brightness settings, Notifications, Keyboard, and Power options.

## ③ Windows key

The Windows key toggles between the desktop and the new Start screen. For the most part, you'll probably use this key to get back to the desktop because there really isn't much you're going to do with the Start screen on a server. If your keyboard doesn't have a Windows key, you can also use the Ctrl + Esc key combination.

## ④ Ctrl+C: copy; Ctrl+X: cut; Ctrl+V: paste; Ctrl+Z: undo

This venerable set of keyboard shortcuts is still quite viable in Windows Server 2012. You can use them in Notepad and Windows Explorer, as well as other applications such as Microsoft Office.

## ⑤ Alt+F4

The Alt + F4 keyboard shortcut closes a program, just like clicking the X icon in the upper-right corner of desktop applications. The Alt + F4 keyboard combination also works for newer Windows 8-style applications, which is useful if you can't figure out how to get out of one of those apps.

## ⑥ Windows key+E

The Win + E keyboard shortcut opens Windows Explorer. Win + E works on both the desktop and the Start screen. When you use this keyboard shortcut, it opens Windows Explorer to the Computer folder.

## ⑦ Windows key+Shift+Tab

By this time, I'm sure you're beginning to see how important keyboard shortcuts are to Windows Server 2012. The Win + Shift + Tab keyboard combination cycles through open applications. This combination also works with [Windows 8](#) applications.

## ⑧ Windows key+R

The Win + R keyboard combination opens a Run prompt where you can enter the path and name of a program to run. The program launches with the same privileges that you're logged in with. If you're logged in as Administrator, the program runs with administrative privileges.

---

**The Windows key toggles between the desktop and the new Start screen.**

---

## ⑨ Windows key+Q

Another useful keyboard shortcut in Windows Server 2012, Win + Q opens the Apps Search dialog box, which lets you enter the name of the application you want to run. In addition, a listing of all installed applications is shown on the screen.

## ⑩ Windows key+L

Found in earlier versions of Windows Server, the Win + L combination locks the computer and requires a login to re-access the system. Win + L is particularly useful if you need to step away from your system for a few minutes. ■



# Businesses of all Sizes Improve Customer Relations and Reduce Costs with Email2DB

Regardless of business size or market, competent communication with customers and leads is a must. Without the ability to quickly and capably respond to inquiries, businesses struggle to demonstrate the supportive backing that customers demand. Ironically, with so many communications channels and mediums available today, too many organizations rely upon a haphazard communications infrastructure that overly burdens employees and commonly fails to effectively demonstrate organizational priorities.

To address these concerns, Parker Software created Email2DB to provide easy parsing, extraction, and customizable responses to messaging from a wide variety of sources in a single, powerful, flexible, and easy-to-use solution.

## ► Email2DB Success Stories

Xerox used Email2DB to enable better support for its multi-national clients. Without Email2DB, support tickets resulted in large amounts of time involved with administration and technician follow-up. With Email2DB, Xerox was able to simultaneously reduce support costs and free up technicians to better service customer needs and requests.

With Email2DB, Wesleyan Assurance Society implemented a highly automated solution

*Parker Software develops software solutions designed to help businesses be more successful online. WhosOn is a Live Chat solution, helping to improve online conversion and customer support through improved customer engagement and real-time tracking. Email2DB is a business process automation solution aimed at improving business efficiency through parsing incoming messages.*

that enables existing consumers of its financial products and services to learn about new services, offerings, and benefits. By coordinating user interactions from a self-service website, telephone calls, or interactive chats against automated email communications facilitated by Email2DB, Wesleyan is able to better service customer needs interactively. Best of all, Email2DB tracks customer exploration and inquiries against automated responses to simultaneously route customers to specialized customer account representatives and provide those representatives with context and information about what customers have already learned in order to better meet their needs.

## ► How Email2DB Works

Email2DB users set up accounts—or specific content types to evaluate and parse—including databases, emails, RSS feeds, tweets,

SharePoint, files, web pages, and others. Users then define processing rules specifying execution frequencies and definitions for the specific content to parse, evaluate, and extract—along with triggers that define exactly how to respond to particular values and content retrieved during evaluation. Triggers, in turn, can be used to send emails, SMS, log to databases, post values to web sites, fire off scripts, or launch applications. Rules, evaluations, filters, processing, and responses can be as simple or as complex as needed and include built-in support for the use of variables and specialized scripting when and where needed—all from a simple, and user-friendly, interface that makes streamlined communications easy to manage.

### ► **Seamless Integration with Microsoft Solutions**

While Email2DB is designed to support a wide variety of sources and destinations, one of its primary strengths is its strong, out-of-the-box integration with Microsoft products and solutions. In addition to native read and write capabilities against traditional data stores such as SQL Server, Excel, and Access, Email2DB also provides seamless integration with Exchange and Microsoft's Dynamics CRM. Email2DB also provides extensive support for interacting with SharePoint, including the ability to interact with SharePoint document libraries, upload files or documents to Lists, or read from Lists and even update SharePoint Sites.

### ► **A Sample Use Case—Email2DB in Action**

Suppose a startup creates a software solution

so compelling that high percentages of trial-version users purchase the solution. To help bolster sales, this startup will want to place as many trial versions as possible, using every medium and outlet available. With Email2DB, it is trivial to respond to email, web inquiries, and even twitter requests for trial software licenses. As such, by parsing inbound requests to extract pertinent information, Email2DB can simultaneously create a catalog of potential customers while instantly returning trial licenses to all inbound requests 24x7x365. Email2DB can then regularly query for users who have activated their trial licenses within the last few days and send follow-up links to tutorials, FAQs, and other resources to help familiarize users with the product's benefits to improve conversion. Then, for trial users who haven't purchased within a specified period of time, Email2DB can send promotional discounts to help incentivize conversion.

### ► **The Benefits of Email2DB**

Email2DB is flexible, highly configurable, and easy to use. With Email2DB, organizations can streamline communications and responsiveness without getting mired down in costly and complex custom development efforts and can, instead, leverage a single, easy-to-use, and simple-to-manage solution that can improve customer interactions, increase efficiency, and decrease operating costs.

Learn more about how Email2DB has [helped businesses](#) of all sizes, or take a [quick tour](#) of how Email2DB works. ●

# Catching Up with the Cloud Security Alliance

An interview with John Howie  
about cloud security



## Sean Deuby

is technical director for *Windows IT Pro* and *SQL Server Pro* and former technical lead of Intel's core directory services team. He's been a directory services MVP since 2004.

Email



Twitter



Long ago, a [Dilbert cartoon](#) spoke to me in just the right way. The humor of the cartoon (management's ability to decree that a hugely complicated project can be completed by some arbitrary date) really hit home. Of course, this cartoon is still relevant today. One of the challenges IT pros face in this disruptive time of cloud computing is determining a logical and organized way to get started when management tells you they want to "move to the cloud" (whatever that means) in six minutes.

Fortunately there's an organization designed specifically to rescue you: the [Cloud Security Alliance](#). The CSA is "a not-for-profit organization with a mission to promote the use of best practices for providing security assurance within Cloud Computing, and to provide education on the uses of Cloud Computing to help secure all other forms of computing. The CSA is led by a broad coalition of industry practitioners, corporations, associations, and other key stakeholders."

At the [RSA Conference 2013](#), I was able to talk with John Howie, the CSA's chief operating officer (and longtime contributor to *Windows IT Pro*), and get his take on cloud security and what the CSA offers to the IT pro.

**Sean Deuby:** How did the CSA get started?

**John Howie:** Four or five years ago, a group of professional colleagues, including current Executive Director Jim Reavis, realized there was a

lack of information about how to move to the cloud while maintaining data security and privacy. This lack of knowledge was inhibiting cloud computing adoption. The CSA was formally announced at the RSA conference as a community where security-minded professionals could get together and share knowledge. The membership quickly exploded because a centralized, vendor-agnostic clearing house was badly needed. Every cloud service provider (CSP) had its own set of documentation, but it was all in different places on the CSP's own websites and stored in different formats. The CSA worked with the CSPs to take their best practices and put them in a guidelines document for all to use.

**Deuby:** What areas does the CSA cover?

**Howie:** Beginning with the original “[Security Guidance for Critical Areas of Focus in Cloud Computing](#)” document (now in its third version), the scope of the CSA's documents grows as new areas (e.g., [mobile computing](#)) grow to a critical mass of interest. Another example is [security as a service](#) (SecaaS)—a huge project for the CSA—and how security has evolved in a hybrid computing environment. We've also opened a [Legal Information Center](#), where IT pros can ask general legal questions about cloud computing (e.g., data restrictions in European countries) and get clear, jargon-free answers.

**Deuby:** What new initiatives do you have going on?

**Howie:** It's not commonly known that major CSPs don't currently provide anything more than the very highest-level availability reports for their services. The CSA is working on building protocols and a framework to eventually allow cloud service consumers to continuously monitor major CSPs at a daily, weekly, or monthly level. We also have the [Security, Trust, and Reliability \(STAR\) initiative](#), which is a place where CSPs can upload their security documentation into



the CSA's framework. This initiative allows a potential customer to easily compare vendors on an apples-to-apples basis.

**Deuby:** How does the IT pro start using the CSA?

**Howie:** Download and read the [Security Guidance document](#). It really is the launching point into everything we do. It covers several areas, and with this document you can gain a high-level understanding of the critical areas of focus for cloud computing. The next step is to read our document "[The Notorious Nine: Cloud Computing Top Threats in 2013](#)." This is a good way to look at a generic set of concerns when you're considering moving to the cloud.

**Deuby:** Are there any other resources that IT pros should be looking at in this area?

**Howie:** Microsoft has released the [Microsoft Cloud Readiness Tool](#), which is a questionnaire-based system that measures the maturity of an organization and its readiness for cloud computing. When you use the tool, it provides you with a results scorecard and guidance about areas you need to focus on. The questions in the questionnaire are mapped back to CSA guidance, and the tool is endorsed by CSA. The results work for any CSP, not just Microsoft.

## How to Get Involved

The CSA is a growing organization, and one you should look into. All the intellectual property is free for personal use, commercial use, or commercial re-use. You can also join the CSA on [LinkedIn](#) to join the many discussions. Individual membership is free; the alliance is funded by its [impressive list of almost 150 corporate members](#). Howie states the CSA mission very simply: "Our only goal is to advance the state of cloud computing security." ■

# The Claims Rule Language in AD FS

**M**icrosoft Active Directory Federation Services (AD FS) uses the Claims Rule Language to issue and transform claims between claims providers and relying parties. Dynamic Access Control, introduced with [Windows Server 2012](#), also uses this common language. The flow of claims follows a basic pipeline. The rules we create define which claims are accepted, processed, and eventually sent to the relying party. In this article, I'll go over the basics of how AD FS builds claims, then dive deep into the language that makes it all work. At the end, you should be able to read a claim rule, understand its function, and write custom rules.

## The Basics

Before diving into the language used to manipulate and issue claims, it's important to understand the basics. A *claim* is information about a user from a trusted source. The *trusted source* is asserting that the information is true, and that source has authenticated the user in some manner. The *claims provider* is the source of the claim. This can be information pulled from an attribute store such as Active Directory (AD), or it can be a partner's federation service. The *relying party* is the destination for the claims. This can be an application such as Microsoft SharePoint or another partner's federation service.

A simple scenario would be AD FS authenticating the user, pulling attributes about the user from AD, and directing the user to an application to consume. The scenario can be more complex by adding partner federation services. In any scenario, we're taking information from some location and sending it somewhere else. Figure 1 shows a sample relationship between federation servers and an application.



**Joji  
Oshima**

is a senior support escalation engineer in Windows Commercial Technical Support focusing on authentication and AD FS. He is a contributor to the Ask the Directory Services Team blog and the [TechNet Wiki](#).




**Email**



**Blog**

**Figure 1**  
Sample Relationship  
Between Federation  
Servers and an  
Application



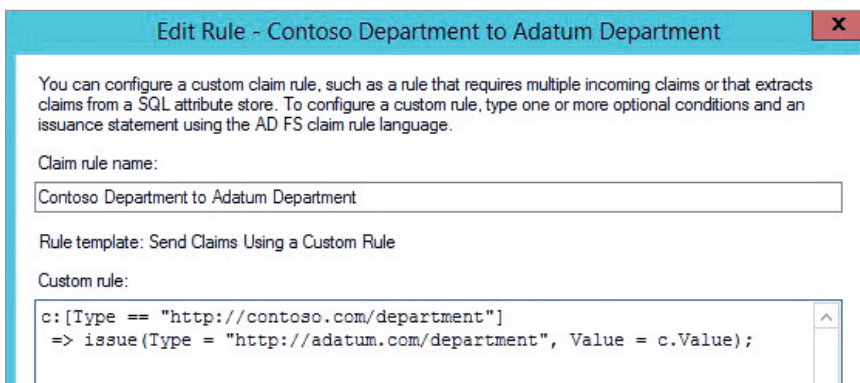
## Claim Sets

You need to understand claim sets in relation to the claims pipeline. When claims come in, they're part of the incoming claim set. The claims engine is responsible for processing each claim rule. It examines the incoming claim set for possible matches and issues claims as necessary. Each issued claim becomes part of the outgoing claim set. Because we have claim rules for claims providers and relying parties, there are claim sets associated with each of them.

1. Claims come in to the claims provider trust as the incoming claim set.
2. Claim rules are processed, and the output becomes part of the outgoing claim set.
3. The outgoing claim set moves to the respective relying party trust and becomes the incoming claim set for the relying party.
4. Claim rules are processed, and the output becomes part of the outgoing claim set.

## General Syntax of the Claims Rule Language

A claim rule consists of two parts: a condition statement and an issuance statement. If the condition statement evaluates true, the issuance statement will execute. The sample claim rule that Figure 2 shows takes an incoming Contoso department claim and issues an Adatum department claim with the same value. These claim types are uniform resource identifiers (URIs) in the HTTP format. URIs aren't URLs and don't need to be pages that are accessible on the Internet.



**Figure 2**  
A Simple Claim Rule

## Condition Statements

When a rule fires, the claims engine evaluates all data in the incoming claim set against the condition statement. Any property of the claim can be used in the condition statement, but the most common are the claim type and the claim value. The format of the condition statement is *c:[query]*, where the variable *c* represents a claim currently in the incoming claim set. The simple condition statement

```
c:[type == "http://contoso.com/department"]
```

checks for an incoming claim with the claim type *http://contoso.com/department*, and the condition statement

```
c:[type == "http://contoso.com/department", value == "sales"]
```

checks for an incoming claim with the claim type *http://contoso.com/department* with the value of *sales*. Condition statements are optional. If you know you want to issue a claim to everyone, you can simply create a rule with the issuance statement.

## Issuance Statements

There are two types of issuance statements. The first is ADD, which adds the claim to the incoming claim set, but not the outgoing set. A

typical use for ADD is to store data that will be pulled in subsequent claim rules. The second is ISSUE, which adds the claim to the incoming and outgoing claim sets. The ISSUE example

```
=> issue(type = "http://contoso.com/department", value =
    "marketing");
```

issues a claim with the type *http://contoso.com/department* with the value of *marketing*. The ADD example

```
=> add(type = "http://contoso.com/partner", value = "adatum");
```

adds a claim with the type *http://contoso.com/partner* with the value of *adatum*. The issuance statement can pull information from the claim found in the condition statement, or it can use static information. The static data example

```
c:[type == "http://contoso.com/emailaddress"]
=> issue(type = "http://contoso.com/role", value = "Exchange User");
```

checks for an incoming claim type *http://contoso.com/emailaddress* and, if it finds it, issues a claim *http://contoso.com/role* with the value of *Exchange User*. The static data example

```
c:[type == "http://contoso.com/role"]
=> issue(claim = c);
```

checks for an incoming claim type *http://contoso.com/role* and, if it finds it, issues the exact same claim to the outgoing claim set. An example of pulling data from the claim,

```
c:[type == "http://contoso.com/role"]
=> issue(type = "http://adatum.com/role", value = c.Value);
```

checks for an incoming claim type *http://contoso.com/role* and, if it finds it, issues the exact same claim to the outgoing claim set.

## Multiple Conditions

Another possibility is to use multiple conditions in the condition statement. The issuance statement will fire only if all conditions are met. Each separate condition is joined with the `&&` operator. For example,

```
c1:[type == "http://contoso.com/role", value=="Editor"] &&
c2:[type == "http://contoso.com/role", value=="Manager"]
=> issue(type = "http://contoso.com/role", value = "Managing Editor");
```

checks for an incoming claim with the type *http://contoso.com/role* with a value of *Editor* and another incoming claim with the type *http://contoso.com/role* with a value of *Manager*. If the claims engine finds both, it will issue a claim with the type *http://contoso.com/role* and the value of *Managing Editor*.

The values of claims in any condition can be accessed and joined using the `+` operator. For example,

```
c1:[type == "http://contoso.com/location"] &&
c2:[type == "http://contoso.com/role"]
=> issue(type = "http://contoso.com/targetedrole", value = c1.Value
+ " " c2.Value);
```

checks for an incoming claim of type *http://contoso.com/location* and a separate incoming claim with the type *http://contoso.com/role*. If it finds both, it will issue a claim with the type *http://contoso.com/targetedrole*, combining the values of the incoming roles.

## Aggregate Functions

Up to this point, each claim rule checks individual claims or groups of claims and fires each time there's a match. There are some circumstances



in which this behavior isn't ideal, however. For example, you might want to look at the entire incoming claim set and make a condition statement based on that. In such cases, you can use the EXISTS, NOT EXISTS, and COUNT functions. The EXISTS function checks whether there are any incoming claims that match; if there are, it fires a rule. The NOT EXISTS function checks whether there are any incoming claims that match; if there aren't, it fires a rule. The COUNT function counts the number of matches in the incoming claim set.

The EXISTS example

```
EXISTS([type == "http://contoso.com/emailaddress"])
=> issue(type = "http://contoso.com/role", value = "Exchange User");
```

checks for any incoming claims with the type *http://contoso.com/emailaddress* and, if it finds any, issues a single claim with the type *http://contoso.com/role* and the value of *Exchange User*. The NOT EXISTS example

```
NOT EXISTS([type == "http://contoso.com/location"])
=> add(type = "http://contoso.com/location", value = "Unknown");
```

checks for any incoming claims with the type *http://contoso.com/location* and, if it doesn't find any, adds a single claim with the type *http://contoso.com/location* with the value of *Unknown*. The COUNT example

```
COUNT([type == "http://contoso.com/proxyAddresses"]) >= 2
=> issue(type = "http://contoso.com/MultipleEmails", value
    = "True");
```

checks for any incoming claims with the type *http://contoso.com/proxyAddresses* and, if there are two or more, issues a single claim with the type *http://contoso.com/MultipleEmails* with the value of *True*.

## Querying Attribute Stores

By default, AD is the only attribute store created when you install AD FS. You can query LDAP servers or SQL Server systems to pull data to be used in a claim. To utilize another attribute store, you first create the attribute store and enter the appropriate connection string. Figure 3 shows how to create an LDAP server as an attribute store.



**Figure 3**  
Creating an LDAP  
Server as an Attribute  
Store

Once you create the store, you can query the store from a claim rule. For an LDAP attribute store, the query should be in this format:

```
query = <query_filter>;<attributes>
```

The parameter sent into the query is represented with the {0} operator. If multiple parameters are sent, they would be {1}, {2}, and so on. For example,

```
c:[Type == "http://contoso.com/emailaddress"]
=> issue(
    store = "LDAP STORE",
    types = ("http://contoso.com/attribute1", "http://contoso.com/attribute2"),
    query = "mail={0};attribute1;attribute2",
    param = c.Value
);
```

queries LDAP STORE for attribute1 and attribute2, where the email address matches, and issues two claims based on the data returned from the query.

A SQL Server attribute store uses the same basic format of the Claims Rule Language; only the query syntax is different. It follows the standard T-SQL format, and the {0} operator is used to pass the parameter. For example,

```
c:[Type == "http://contoso.com/emailaddress"]
=> issue(
    store = "SQL STORE",
    types = ("http://contoso.com/attribute1", "http://contoso.com/attribute2"),
    query = "SELECT attribute1,attribute2 FROM users WHERE email = {0}",
    param = c.Value
);
```

queries SQL STORE for attribute1 and attribute2, where the email address matches, and issues two claims based on the data returned from the query.

## Regular Expressions

The use of regular expressions (Regex) lets you search or manipulate data strings in powerful ways to get a desired result. Without Regex, any comparisons or replacements must be an exact match. This is sufficient for many situations, but if you need to search or replace based on a pattern, you can use Regex. Regex uses pattern matching to search inside strings with great precision. You can also use it to manipulate the data inside the claims.

To perform a pattern match, you can change the double equals operator (==) to =~ and use special metacharacters in the condition statement. If you're unfamiliar with Regex, let's start with some of the common metacharacters and see what the result is when using them. Table 1 shows basic Regex metacharacters and their functions.

Table 1: Basic RegEx Metacharacters and Their Functions

| Symbol | Operation                                    | Example Rule  |
|--------|--|---|
| ^      | Match the beginning of a string              | c:[type == "http://contoso.com/role", Value =~ "^director"]<br>=> issue (claim = c)<br>Pass through any role claims that start with "director"  |
| \$     | Match the end of a string                    | c:[type == "http://contoso.com/email", Value =~ "contoso.com\$"]<br>=> issue (claim = c)<br>Pass through any email claims that end with "contoso.com"   |
|        | OR   | c:[type == "http://contoso.com/role", Value =~<br>"^(director manager)"]<br>=> issue (claim = c)<br>Pass through any role claims that start with "director" or "manager"                            |
| (?i)   | Not case sensitive                           | c:[type == "http://contoso.com/role", Value =~ "(?i)^director"]<br>=> issue (claim = c)<br>Pass through any role claims that start with "director," regardless of case                              |
| x.*y   | "x" followed by "y"                          | c:[type == "http://contoso.com/role", Value =~ "(?i)<br>Seattle.*Manager"]<br>=> issue (claim = c)<br>Pass through any role claims that contain "Seattle" followed by "Manager," regardless of case |
| +      | Match preceding character one or more times  | c:[type == "http://contoso.com/employeeId", Value =~ "^0+"]<br>=> issue (claim = c)<br>Pass through any employeeId claims that start with at least one "0"  |
| *      | Match preceding character zero or more times | Similar to above, more useful in RegExReplace() scenarios   |

## RegExReplace

You can also use RegEx pattern matching in replacement scenarios. This is similar to a find-and-replace algorithm found in many text editors, but it uses pattern matching instead of exact values. To use this in a claim rule, use the RegExReplace() function in the value section of the issuance statement.

The RegExReplace function accepts three parameters:

- The first is the string in which you're searching. You'll typically want to search the value of the incoming claim (c.Value), but this could be a combination of values (c1.Value + c2.Value).

- The second is the RegEx pattern you're searching for in the first parameter.
- The third is the string value that will replace any matches found.

The RegExReplace example

```
c:[type == "http://contoso.com/role"]
=> issue (Type = "http://contoso.com/role", Value =
    RegExReplace(c.Value, "(?i)director", "Manager");
```

passes through any role claims. If any of the claims contain the word *Director*, RegExReplace will change it to *Manager*. For example, *Director of Finance* would pass through as *Manager of Finance*.

If you combine the power of RegEx pattern matching with the concepts mentioned earlier in the article, you can accomplish many tasks using the Claims Rule Language.

## Coding Custom Attribute Stores

AD FS gives you the ability to plug in a custom attribute store if the built-in functionality isn't sufficient to accomplish your goals. You can use standard .NET code such as `ToUpper()` and `ToLower()` or pull data from any source through the code. This code should be a class library and will need references to the `Microsoft.IdentityModel` and `Microsoft.IdentityServer.ClaimsPolicy` assemblies.

## Try Custom!

Creating custom rules with the Claims Rule Language gives you more flexibility with claims issuance and transformation. It can take a while to familiarize yourself with the syntax, but it becomes much easier with practice. If you want to dive into this language, try writing custom rules instead of using the templates next time. ■

# Windows 8 Security

## New features and changes

During the development of Windows 8 and following the OS's final release in fourth quarter 2012, the media focused largely on the new Windows 8 Modern UI (previously known as Metro), which was primarily designed for touch-enabled PCs, laptops, and tablets. Less attention has been devoted to Windows 8 security, but there's much more on the table than you might have imagined.

### Windows 8 UI Applications

Microsoft sees the future in Windows 8 applications—and the new programming model that enables developers to quickly provision apps. This model is based on JavaScript with HTML5 and CSS3 as the presentation layer (or alternatively, Visual Basic, C++, or C# and Extensible Application Markup Language—XAML).

This new development environment has given Microsoft the ability to integrate security from the get-go as opposed to bolting on security measures, as we see in traditional desktop applications. As such, low-level changes in Windows 8 are designed to make Modern UI apps more secure than their desktop counterparts.



### Russell Smith

is an independent IT consultant specializing in systems management and security, and author of *Least Privilege Security for Windows 7, Vista and XP* (Packt).



Email



Twitter



## Application Sandboxing

Prior to Windows Vista, access control lists (ACLs) determined which user or users could access files and other system objects, such as registry entries. Integrity levels were introduced in Windows Vista, partly to help implement Internet Explorer (IE) Protected Mode, which restricts websites from modifying objects that have medium or high integrity. Integrity levels add another dimension to Windows security, so that processes and objects that are accessible to a user can also be marked with a level of trust. For example, applications or files that are deemed to be a risk to system security, such as IE or files downloaded from the Internet, run with low integrity and can't modify objects that are marked with a higher integrity level.

AppContainer is a new security mechanism for Windows 8 Modern UI applications. AppContainer is enabled by a new integrity level that blocks Read and Write access to items with higher integrity. This approach differs from Windows Vista and Windows 7, in which applications that run with low integrity can read objects with medium or high integrity. All Windows 8 Modern UI applications run with the AppContainer integrity level, except for Internet Explorer 10 (IE10), which runs with medium integrity.

When users log on to Windows 7, each user session is assigned a separate kernel namespace to prevent conflicts. The Windows kernel uses namespaces to hierarchically organize Windows resources (or objects). In Windows 8, apps with the AppContainer integrity level create named kernel objects in a separate namespace from the user session. Unlike other integrity levels in Windows 8, AppContainers can be fine-tuned to suit individual applications. A set of defined capabilities allow the apps to access areas of the OS that are denied by default.

## Declarations

By default, Windows 8 Modern UI applications can access only their own storage areas. To get access to a user library, the network, or a hardware device, an app must declare this intention in its manifest

file. When you install an app from the Windows Store, the declarations about the kind of access the app needs to run are displayed. The user can decide whether to allow this access. These 10 capabilities can be declared:

- `internetClient`
- `internetClientServer`
- `privateNetworkClientServer`
- `documentsLibrary`
- `picturesLibrary`
- `videosLibrary`
- `musicLibrary`
- `enterpriseAuthentication`
- `sharedUserCertificates`
- `removableStorage`

The `enterpriseAuthentication` capability allows apps to impersonate the credentials of the logged-in user on the network. All the other capabilities are self-explanatory, but note that apps can access files on removable storage only when the file type is specifically declared in the manifest. Files on HomeGroup network shares aren't accessible. At the time of writing, Windows 8 Modern UI apps can't access SQL Server for local data storage, but Microsoft is expected to introduce more capabilities in future Windows releases. For more information on declaring capabilities for Modern UI applications, see the [Capabilities topic in the Windows Dev Center](#).

## IE Enhanced Protected Mode

Tabs in the desktop version of IE10 also run in an AppContainer sandbox when Enhanced Protected Mode (EPM) is enabled. Protected Mode uses Mandatory Integrity Control to help prevent the installation of malicious code or changes to system files if the browser is exploited. EPM builds on Protected Mode by adding a series of new security features.

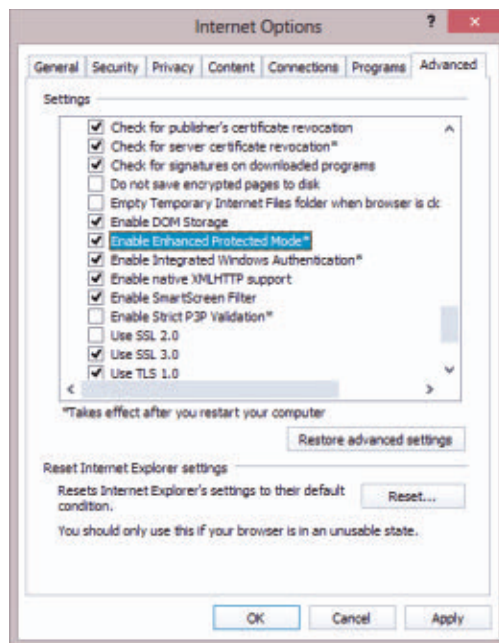
If you've loaded Windows 8 and started IE, you might have noticed that IE10 runs in mixed 32-bit/64-bit mode by default. The window borders and menus run in 64-bit processes and the tabs run in 32-bit processes.

Whenever possible, run 64-bit processes to take advantage of advanced security features, such as Address Space Layout Randomization (ASLR)—more on ASLR later in this article. Some toolbars and plug-ins are incompatible with 64-bit IE, so EPM is disabled by default in desktop IE. Because the Windows 8 Modern UI app version of IE has a no-plug-in architecture, EPM is unlikely to cause any compatibility problems and is enabled by default. If desktop IE comes across a plug-in that's incompatible with EPM, then the plug-in is automatically disabled. You can then choose whether to disable EPM specifically for the website in question.

To give extra protection to personal data, EPM uses a broker process when a user tries to upload a document to a website. This process gives IE access to the document only when the user clicks Open

in the File Upload dialog box. When EPM is enabled, IE has access to user documents only when necessary, rather than all the time. Tabs are also restricted from accessing local sites running on the corporate Intranet and from accessing users' domain-login credentials. Finally, EPM tabs can't run as a local webserver. EPM can be switched on from the Advanced tab in the Internet Settings applet in Control Panel or by using Group Policy, as Figure 1 shows.

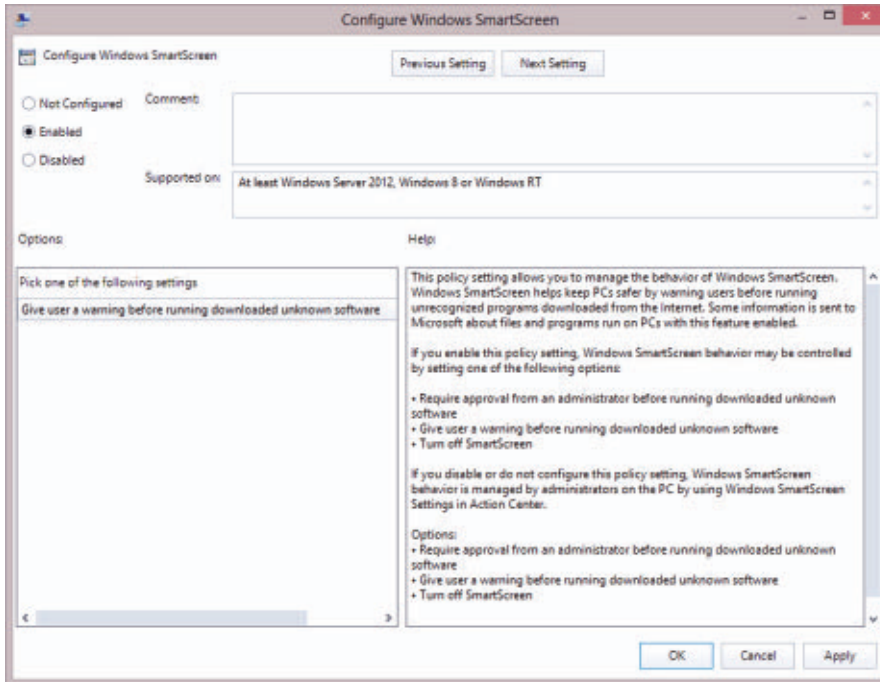
**Figure 1**  
Enabling EPM in IE10



## SmartScreen

First introduced to check the integrity of websites that are loaded in IE, the SmartScreen filter has been expanded in Windows 8. The filter now also checks all files that are downloaded from the Internet against the Microsoft reputation service.

In Windows 8, SmartScreen is configured by default to require administrator approval before a file with no reputation can be run. You can change SmartScreen settings in the Action Center or via Group Policy under Computer Configuration, Administrative Templates, Windows Components, File Explorer, as Figure 2 shows. The SmartScreen filter can also be controlled via Group Policy for IE security zones.



**Figure 2**  
Configuring  
SmartScreen for File  
Explorer in Windows 8

## Platform Integrity Architecture: Measured and Secured Boot

Hardware-based security has been beefed up in Windows 8. It now supports several new features that are enabled by Unified Extensible Firmware Interface (UEFI) firmware and the Trusted Platform Module

(TPM), which can also be used by BitLocker in Windows Vista or later. The Measured Boot and Secure Boot protocols help identify rootkits and prevent malicious code from hiding from the OS. Measured Boot requires TPM but will work with conventional BIOS.

**Secure Boot.** Windows 8 supports the UEFI Secure Boot protocol, as part of the Windows 8 secured boot architecture. Secure Boot is designed to work with hardware devices that run a UEFI-compatible BIOS, to verify the integrity of the pre-Windows environment and stop malware from modifying the firmware and injecting itself pre-boot. An independent forum is developing UEFI to use the capabilities of modern hardware so that the pre-OS environment can communicate directly with the hardware, using fast block I/O instead of legacy software interrupts. The Trusted Computing Group (TCG) defines the [UEFI specifications](#).

Secure Boot prevents the device firmware from booting malicious OS loaders, which can remain undetected by Windows. Windows 8 platform integrity architecture also includes Early Launch Anti-Malware (ELAM) and the ability to verify client health by comparing the system start state, as recorded in the device's TPM, with data held in a remote verifier.

Public Key Infrastructure (PKI) is used to establish a chain of trust. During manufacture, OEMs install a platform key that protects firmware from malicious changes. When a computer with UEFI Secure Boot starts, the firmware verifies the state of boot-loader files and the firmware on devices such as network and video cards, before handoff. Verification occurs by comparing the signatures stored in the various firmwares against databases of allowed and disallowed signatures. Your motherboard must support UEFI, which must be enabled in the firmware before Windows 8 is installed. After UEFI is enabled, Windows 8 can be installed from DVD media or from a bootable FAT32 USB stick. UEFI requires that Windows be installed on a GUID Partition Table (GPT) disk partition; if you have a previous installation of Windows on a Master Boot Record (MBR) partition,

you'll need to wipe it and convert to GPT. You can do this by using the Diskpart command from Windows Preinstallation Environment (Windows PE), if necessary.

After the boot partition is prepared, you can start installing Windows 8. To install Windows 8 in UEFI mode and not BIOS mode, make sure that you boot from a designated Extensible Firmware Interface (EFI) drive when you boot from DVD or USB to install Windows. The simplest way to ensure this is to set the boot drive specifically in the firmware. You'll then be able to see drives marked as EFI.

UEFI Secure Boot requires four partitions, but the install process does all the necessary configuration for you. Table 1 shows the required partitions.

Table 1: UEFI Default Disk Partitions

| Partition                               | Size                 |
|---|----------------------|
| Windows Recovery Environment (RE) tools | 300MB                |
| System                                  | 100MB                |
| Microsoft Reserved Partition (MSR)      | 128MB                |
| Windows                                 | Remaining disk space |

After Windows 8 is installed, you can check that UEFI Secure Boot is enabled by running the following PowerShell command:

```
confirm-SecureBootUEFI
```

The command returns a value of True if Windows 8 is running in UEFI mode.

**Measured Boot.** Measured Boot is a new feature, available to anti-virus software and building on UEFI's Secure Boot. It confirms the health of a machine by using a log file, stored on the TPM chip, that's generated every time the system boots. This log contains a list



of hashes that are used to confirm the integrity of the drivers and components that are loaded during the boot process, before antivirus is started. The log file can also be protected by a cryptographic key that's issued to the TPM. When fully loaded, antivirus software running in Windows can inspect the hashes to check for any unauthorized changes to boot components and drivers.

Remote attestation allows the antivirus software running on the client to send the Measured Boot log file to a server for verification. In this way, the client PC doesn't rely on itself to determine its own health. Servers are deemed to have a higher level of assurance than PCs, so confirming the claims made by client PCs in this way is preferable.

## ELAM System

If enabled, the Windows ELAM system driver is the first to load. It allows an ELAM driver, provided by antivirus software, to categorize drivers as good or bad (see Table 2 for classifications). This information is then passed back to the Windows ELAM driver, which determines whether successive boot drivers should be initialized, based on a driver initialization policy. The ELAM Boot-Start Driver Initialization Policy can be found in Group Policy under Computer Configuration, Administrative Templates, System, Early Launch Anti-Malware.

Table 2: ELAM Boot-Start Driver Classification

| Classification             | Description  |
|----------------------------|--|
| Good                       | The driver has been signed and has not been tampered with.   |
| Bad                        | The driver has been identified as malware. You should not allow known bad drivers to be initialized.   |
| Bad, but required for boot | The driver has been identified as malware, but the computer cannot successfully boot without loading this driver.  |
| Unknown                    | This driver has not been attested to by your malware detection application and has not been classified by the Early Launch Anti-Malware Boot-Start Driver. |

When enabled, the Boot-Start Driver Initialization Policy can be set to initialize only good drivers; good and unknown; good, unknown, and bad but critical; or all drivers. Assuming that your hardware has a TPM, Windows Defender supports ELAM out of the box.

## Advanced Exploit Mitigations

ASLR, which has been part of Windows since Windows Vista, randomizes the location of system libraries in memory to make exploits more difficult. ForceASLR in Windows 8 has been improved to randomize the location of DLLs that don't declare themselves compatible with ASLR. [ForceASLR](#) is also available for Windows 7 and Windows Server 2008 R2 as an optional download.

Also new to Windows 8 is High Entropy ASLR (HEASLR), which brings a greater level of randomization to ASLR but requires DLLs to opt in before it's applied. If you run devices with Intel's new Ivy Bridge processor architecture, then Windows 8 uses the new random number generator (Intel Secure Key Technology), providing a greater level of randomization than when relying on the system clock. Windows 8 randomizes some low-level memory-allocation functions, randomizing not only system libraries but also standard Windows applications. IE10 uses the full set of ASLR functions. Windows 8 UI apps and Windows Services always opt in for HEASLR.

Windows 8 can't be installed on hardware that doesn't support an Intel NX (No eXecute) bit or its equivalent. Hardware data execution prevention (DEP), which is designed to prevent applications from executing code in non-executable areas of memory, is a system requirement. The Windows 8 kernel now runs in and allocates non-executable memory, making it more difficult to attack. Supervisor Mode Execution Prevention (SMEP) in Intel Ivy Bridge processors helps prevent attackers from exploiting bugs in the system kernel. SMEP works by not allowing the kernel to execute attack code residing in memory that's allocated to a malicious process. For more information on DEP and ASLR, see "[Vista and Server 2008 Malware Protection Gems](#)."

## Access Control and Data Management

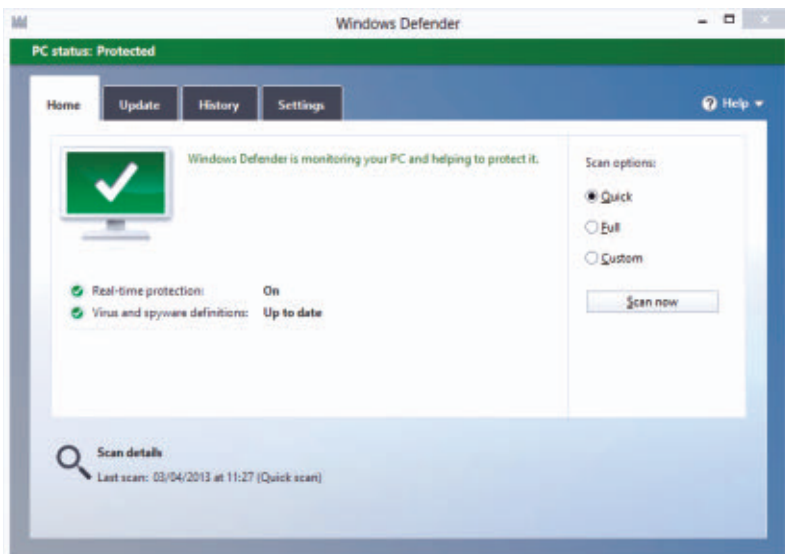
Dynamic Access Control (DAC) is new in [Windows Server 2012](#). It promises to change the way organizations manage access to corporate data. The traditional NTFS-based system of ACLs is too unwieldy to manage on a large scale, partly because it doesn't offer any form of centralized management. Windows 8 is the only Microsoft OS to support claims-based authorization, which is a requirement for working with DAC in Windows Server. Compound claims (user and device claims) also require Windows 8. Changes to DirectAccess and the ability to manage virtual smart cards give Windows 8 users the ability to use the device as a network access token, negating the need for a physical smart card. For more information on DAC, see the Windows Server Blog post "[Introduction to Windows Server 2012 Dynamic Access Control](#)" or go to "[Understanding and Evaluating Virtual Smart Cards](#)" for a white paper on virtual smart cards.

**Figure 3**

Antivirus Built-In to  
Windows Defender

## Windows Defender

Now incorporating functionality that was available separately as Microsoft Security Essentials, Windows 8 computers have antivirus and spyware out of the box via Windows Defender, as Figure 3 shows. Although Windows Defender is likely to appeal mostly to consumers, small businesses, or those who simply don't want to spend money on security software, providing Windows with basic antivirus capabilities can only be good for the overall security landscape. Windows Defender can't be managed centrally, so large



organizations will still need to pay for antivirus to get full management and monitoring functionality.

### Boost Hardware-Based Security and More

Windows 8 offers much more than just a shiny new interface. Some (but not all) security improvements rely on new hardware developed by Microsoft and Intel. Although there's no such thing as 100-percent secure, Windows 8 makes it significantly more costly for an attacker to penetrate systems that are properly configured with security in mind. ■

## John Savill Master Series

*Encompasses eight key technology areas  
24 eLearning sessions  
Beginning July 25th*

*Register now  
[WindowsITPro.com/MasterSeries](http://WindowsITPro.com/MasterSeries)*

# Getting Started with Hyper-V in Windows Server 2012

## Major updates make Hyper-V a must



### John Savill

is a Windows technical specialist, an 11-time MVP, and an MCSE for Private Cloud and Server Infrastructure 2012. He's a senior contributing editor to *Windows IT Pro* and his latest book is *Microsoft Virtualization Secrets* (Wiley).

Email



Twitter



Website



Blog



**W**indows Server 2012 is shaking up the world of server virtualization. In particular, major updates to Hyper-V make the hypervisor a viable option for enterprises that previously might not have considered it.

Windows Server 2012 Hyper-V features the highest levels of scalability available, with 64 virtual CPU (vCPU) virtual machines (VMs) that can have up to 1TB of assigned memory. You can create virtual hard disks (VHDs) of 64TB, thanks to the new VHDX format, which removes the need to use pass-through storage. Put those items together and pretty much any workload can be virtualized with Hyper-V.

Of course, all the scalability in the world doesn't mean much without features that can take advantage of it. Fortunately, Windows Server 2012 delivers those as well. Storage live migration, which enables storage of a VM, lets you move VMs without downtime. Shared-nothing live migration allows VMs to be migrated between Windows Server 2012 Hyper-V hosts with no downtime, even without being clustered and without common storage. These features offer complete mobility of VMs in the data center.

New network and storage options—including single root I/O virtualization (SR-IOV), Server Message Block (SMB) 3.0, virtual Fibre Channel, and network virtualization—make Hyper-V an appealing hypervisor choice. However, most organizations haven't considered Hyper-V before and might not be clear about how to get started. In this article, I cover some of the basics of getting up and running with Windows Server 2012 Hyper-V.

## What Hardware Do I Need?

For a basic, single-server setup, you need a server with a 64-bit processor that supports hardware-assisted virtualization. For Intel processors, this is the Intel Virtualization Technology (VT) feature; for AMD processors, you need AMD virtualization (AMD-V). Pretty much any server processor manufactured within the past 5 years should have this capability. But if you aren't sure about your hardware, download and run the [Coreinfo utility](#), with the -v switch, from an elevated command prompt. This action will show whether the processor supports virtualization and whether it supports Second Level Address Translation (SLAT)—also called Extended Page Tables (EPT) by Intel and AMD.

The output in Figure 1 shows that Intel hardware-assisted virtualization is enabled, which is all we need to get started. SLAT isn't required for Hyper-V to function, but it does improve performance. Therefore, the use of SLAT is preferred when possible and is crucial for virtual-desktop environments such as virtualized Remote Desktop Services servers and virtual desktop infrastructure (VDI) environments.

Windows Server 2012 doesn't have any of the supported virtual processor-to-logical processor ratio limitations that were present in

Figure 1: Sample Coreinfo Output

```
C:\>coreinfo -v
Coreinfo v3.2 - Dump information on system CPU and memory topology
Copyright (C) 2008-2012 Mark Russinovich
Sysinternals - www.sysinternals.com
Intel(R) Core(TM)2 Quad CPU    Q6600  @ 2.40GHz
Intel64 Family 6 Model 15 Stepping 11, GenuineIntel
HYPERVISOR      -      Hypervisor is present
VMX              *      Supports Intel hardware-assisted virtualization
EPT              -      Supports Intel extended page tables (SLAT)
```



previous versions. (An 8:1 ratio was supported in Windows Server 2008 R2 for VMs running server OSs.) Basically, if the server is handling the load to your satisfaction, that's good enough for Microsoft!

The amount of memory required depends completely on the amount that you want to allocate to VMs. I generally carve out around 2GB of memory for the virtualization host, then base additional memory on the amount I need for VMs. For large-scale virtualization environments, servers with 96GB or 192GB of memory are common. But in a lab environment, you need only enough to run your desired virtual load.

Each VM has one or more VHDs. For Windows Server 2012, use the new VHDX format, which not only supports 64TB VHDs (up from 2TB with the old format) but also has been re-architected to offer near bare-metal disk performance. This is true even for dynamic VHDs, which use up only a small amount of disk space initially and grow the file as data is written to the VHD. You also have the option to create a fixed-size VHD. This option is typically used in production environments, both for legacy performance reasons and to avoid the possibility of running out of physical disk space. That's something that can happen as dynamic VHDX files expand, if proper monitoring isn't in place to track the actual physical disk space used. Use the old VHD format only if you need compatibility with older Hyper-V servers, such as Windows Server 2008 R2.

VHDX files can be stored on locally attached storage (i.e., internal disks, in lab environments), although ideally you should use external storage, such as a SAN. New to Windows Server 2012 is the ability to use an SMB 3.0 file share to store and run VMs. Using external storage simplifies the backup of virtual environments. As you increase the number of servers, external storage enables a higher utilization of disk space. The use of a central pool actually makes management easier. External storage is also required if you're going to use Failover Clustering to group multiple hosts into a cluster, allowing VMs to move easily between hosts and automatically restart if a host fails.

(Shared-nothing live migration also allows VMs to move, with no downtime, outside of a cluster, as previously mentioned.) Clusters are also great for maintenance because they allow VMs to be moved to another host (using live migration, which prevents downtime to the VM) while the original host is patched and rebooted, then moved back while the next node is evacuated of VMs, patched, and rebooted. Every host in the cluster can be patched, without any downtime to VMs. Windows Server 2012 actually features one-click patching of an entire cluster, using this process.

## What About Network Connections?

Before I talk about the number of network adapters you need, let's review how to use them:

- First, you must be able to manage the Hyper-V host. Therefore, you need a management connection to communicate over the network.
- Second, the VMs most likely need connectivity to the outside world. Private virtual switches can allow communications between VMs only, and internal virtual switches can allow communications between VMs and the Hyper-V host, but neither provides communications to the outside world. You therefore need a network adapter for VM traffic. In a production environment, you likely have at least two network adapters for VM traffic; you can team them to create a single load-balanced, fault-tolerant connection. The option exists to share the network adapter used for VMs with the management OS; in a lab environment, you could use this solution. But ideally, you should separate the management traffic and the virtual network switch that manages VM traffic. If a problem occurs with the virtual switch, you don't want to lose access to the server.
- Third, you need a method for the hosts in a cluster to communicate for internal purposes, such as various IsAlive messages. Typically, this method is a separate network (although networks

that are used for other purposes can—and will—be used if your cluster network is unavailable). In addition to cluster-heartbeat traffic, the cluster network is also used for cluster shared volume (CSV) traffic. This use allows all the cluster hosts to simultaneously access the same set of NTFS LUNs. The CSV traffic typically consists of only metadata changes. However, in some scenarios all storage traffic for certain hosts uses this network. So when using CSV, you should carve out a separate network for the cluster.

- Fourth, you need a dedicated network to ensure a timely migration of VMs between Hyper-V hosts. So you need to allocate a network for live migration.
- Fifth, if you use iSCSI for storage access, then you need a separate network for iSCSI communication.

This demand for five network connections doesn't take into account the use of multiple network adapters for VM traffic or VM teaming (for load balancing and high availability). Nor does it consider the use of multiple iSCSI network connections or Microsoft Multipath I/O (MPIO) for added fault tolerance.

This scenario assumes that you're using 1Gbps networks. The situation is different if you use 10Gbps. There's no sense in having a dedicated 10Gbps network connection for management traffic or CSV traffic. Production environments with 10Gbps likely have two connections, so team them for fault tolerance and then use Quality of Service (QoS) to reserve enough bandwidth for each traffic type, in case of contention. Microsoft details these recommendations in its [“Hyper-V: Live Migration Network Configuration Guide.”](#)

You can use the same approach for 1Gbps connections. Team your connections and use QoS to ensure bandwidth for different traffic types. Another option: Some new platforms have converged fabrics with huge bandwidth pipes that can be virtually carved up into virtual network and storage adapters.

## Which OS Should I Use?

When you have processor, memory, disk, and network worked out, all you need is an OS. But should you use Windows Server 2012 Standard, Windows Server 2012 Datacenter, or the free Microsoft Hyper-V Server 2012? From a Hyper-V feature perspective, all are identical. All three OSs have the same limits, clustering capabilities, and features. The decision depends entirely on which OS you'll be running on the Hyper-V host.

If it's Windows Server, do you intend to freely move the OS instances (i.e., VMs) between hosts? Hyper-V Server 2012 doesn't include any Windows Server guest OS instance rights. This makes sense—the OS is free—and is a great choice if you aren't running the Windows Server OS on the hypervisor. If you're running a VDI environment with [Windows 8](#) VMs, or if you're running only Linux or UNIX VMs, then use Hyper-V 2012.

When you want to run the Windows Server OS in the VMs, Windows Server 2012 Standard includes the right to run two Windows Server VMs. If I wanted to run four VMs with Windows Server, I could buy two copies of Windows Server 2012 Standard. Note that you can still run other VMs with a non-Windows Server OS on the same servers. There's no VM limitation, just a limit on the number of Windows Server guest OS rights running on Hyper-V, VMware, or anything else.

If you want to run numerous Windows Server OS instances, Windows Server 2012 Datacenter includes the right to run an unlimited number of VMs running the Windows Server OS. Consider the price of the Standard and Datacenter versions for your environment. For example, if you're running six or fewer VMs with Windows Server, it's less expensive to buy multiple copies of Standard than to buy Datacenter. But if you're clustering hosts and want to move the VMs, then you have another consideration: Windows Server licenses are tied to a specific piece of hardware and can be moved between servers only every 90 days.

Let's take the example of a branch office with two clustered virtualization hosts. (Again, we're talking about licensing of Windows, not of Hyper-V, so everything I'm talking about is independent of the hypervisor you use.) On each Hyper-V host, there are typically four VMs running, so I need two copies of Windows Server 2012 Standard for each host. But the hosts are clustered so that I can move the VMs between servers. Maybe as part of patching, I want to move all the VMs to host B (which would then end up running eight VMs), while I patch and reboot host A. I then want to move all eight VMs to host A while I patch and reboot host B. I can't do this unless I wait 90 days between the time I move the VMs to host B and the time I move them to host A. And I'd need to wait another 90 days before moving half the VMs back to host B! Plus, if a host fails, I can't move migrated VMs back to the fixed host for 90 days. If I want to move my VMs around freely, then I need to have enough licenses to cover the high watermark of all the VMs that might ever run on one box—eight VMs. To do that, I need four copies of Standard for each server—at which point it makes more sense to just go with Datacenter.

Companies often misunderstood this comparison, which is an important consideration as you plan your licensing. You'll generally use Windows Server 2012 Standard for physical deployments or lightly virtualized environments. You'll typically use Windows Server 2012 Datacenter for true virtualization environments.

The next consideration is whether or not to use the Server Core or *Server with a GUI* installation (previously known as a full installation). In general, use the Server Core configuration level, which requires less patching and therefore fewer reboots. However, the good news is that with Windows Server 2012, you can change the configuration level at any time, with only a reboot required. So, especially if you're new to Windows Server 2012, install the *Server with a GUI* configuration initially. Perform the configuration, get comfortable, then remove the graphical shell and management tools, run

your server at the Server Core configuration level, and manage it remotely from a Windows 8 system. (You can learn more about the different configuration levels in “[Windows Server 2012 Installation Options](#).”)

## Installing Hyper-V

You’ve installed Windows Server 2012, applied the most recent patches, connected to storage, renamed your network adapters to enable easy identification (and teamed them, if required), and configured IP addresses. The next step is to enable the Hyper-V role. This task can be performed graphically, through Server Manager, using the same process that you use to add any other role or feature. Or you can use Windows PowerShell:

```
Install-WindowsFeature Hyper-V  
-IncludeManagementTools -Restart
```

The benefit of using the Server Manager GUI is that it also prompts you to create a virtual switch on a selected network adapter in the server. The virtual network adapters that you configure on your VMs connect to this switch to access the external network. By default, a virtual network adapter is also created on the host OS so that the OS can use that adapter for VM traffic. If you have a dedicated management network adapter, disable the shared adapter after you complete the installation process.

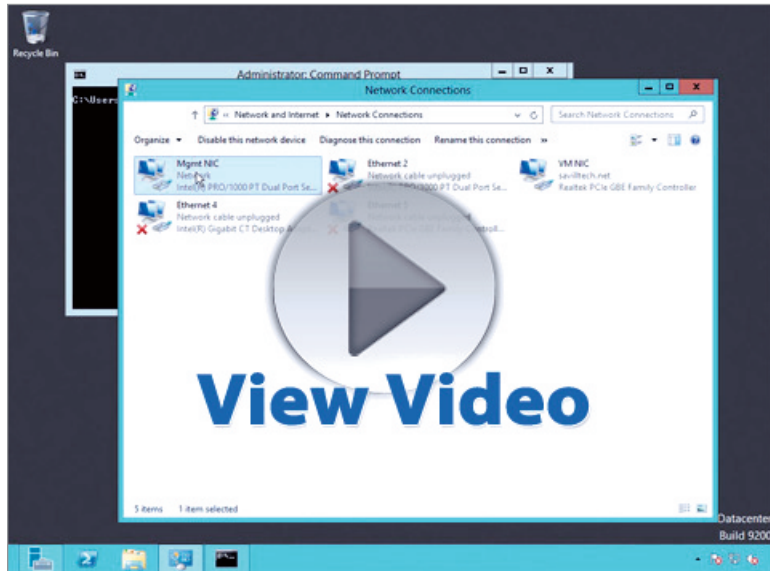
Creating the virtual switch after installation is a straightforward process and can be performed by using PowerShell. The choice of whether to use Server Manager (local or remote) or PowerShell is primarily a matter of preference. If you’re automating the deployment of Hyper-V, then use PowerShell, because manual steps need to be avoided.

In the accompanying video, I quickly walk through the entire Hyper-V installation process, showing the changes to networking.

## Video



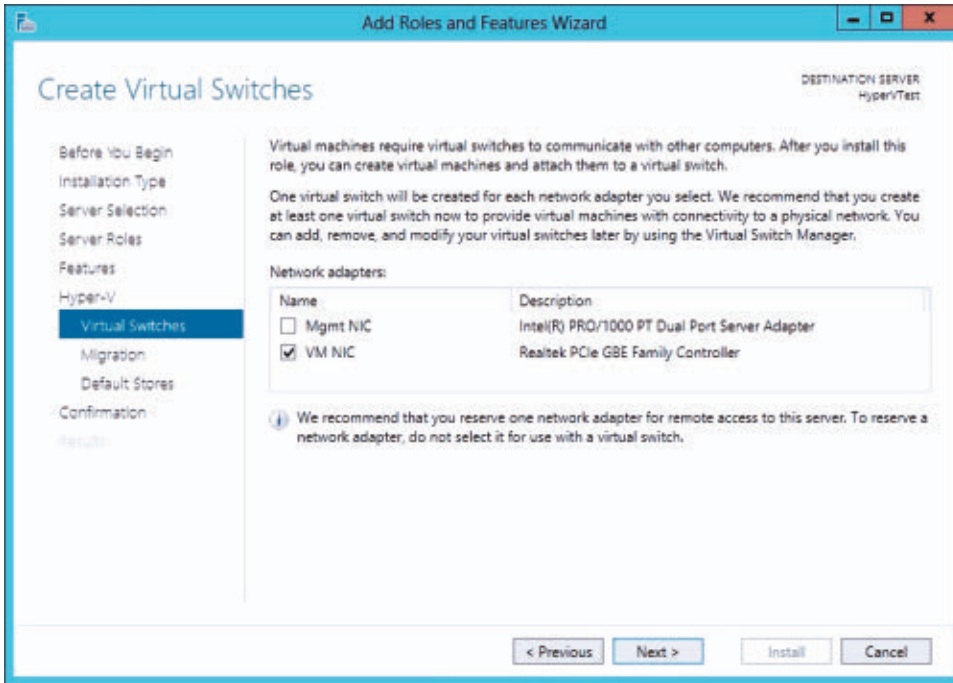
John Savill  
demonstrates the  
Hyper-V installation  
process



The following are the basic steps for using Server Manager:

1. Log on, as an account with administrative credentials, to the server that will be the Hyper-V host and launch Server Manager. Or remotely launch Server Manager with an account that has administrative credentials on the server that will be the Hyper-V host.
2. Select *Add Roles and Features* from the Manage menu.
3. Click Next on the Before You Begin page.
4. On the Installation Type page, choose the Role-based installation type and click Next.
5. On the Server Selection page, from the list of servers in the server pool, choose the server on which to install the Hyper-V role and click Next.
6. Under Server Roles, select Hyper-V and accept the option to automatically install the management tools.
7. On the Create Virtual Switches page, which Figure 2 shows, select the network adapter that you want to use for VM traffic and click Next.





**Figure 2**  
Selecting the Network Adapter to Use for VM Traffic

8. Leave the check box for the option to enable live migrations cleared and click Next. Live migration can easily be added later.
9. Choose new locations for VM storage, or accept the defaults, and click Next.
10. Select the check box to enable automatic restart of the server if required, and click Yes in the displayed confirmation box. Click the Install button.

After the server restarts, you are the proud owner of a Hyper-V virtualization host. Running

```
bcdedit /enum
```

from a command prompt shows that the hypervisor is now autoloading at system startup, as the output in Figure 3 shows.

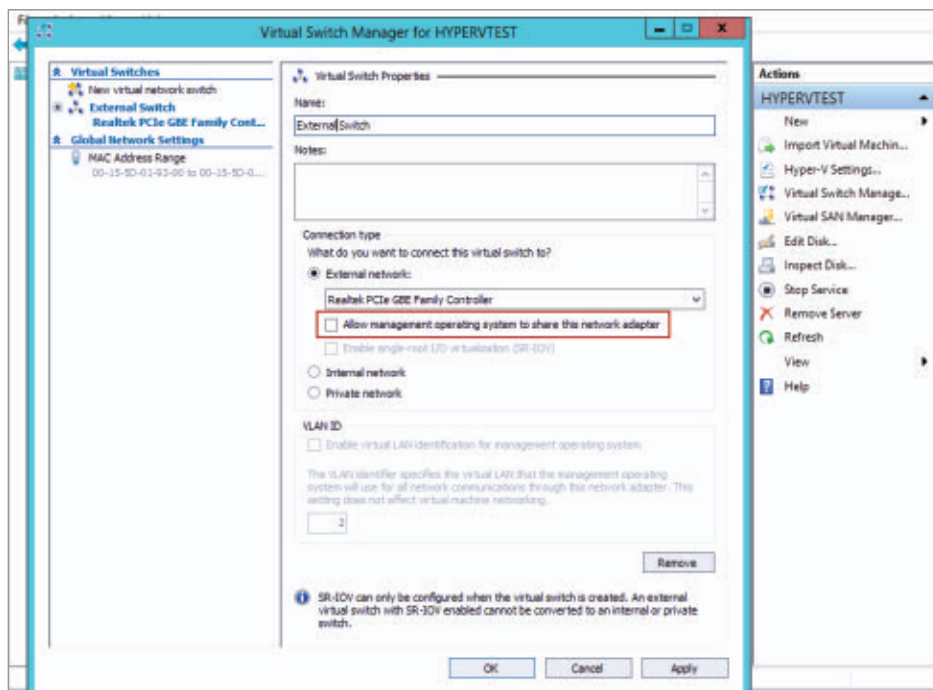
Figure 3: Output Showing Autoloading of Hypervisor at Startup

```

C:\Users\Administrator>bcdedit /enum
Windows Boot Manager
-----
identifier                {bootmgr}
device                    partition=\Device\HarddiskVolume1
description                Windows Boot Manager
locale                    en-US
inherit                    {globalsettings}
bootshutdowndisabled      Yes
default                    {current}
resumeobject               {424ad143-811e-11e2-abd4-ee945fee3b56}
displayorder               {current}
toolsdisplayorder          {memdiag}
timeout                    30
Windows Boot Loader
-----
identifier                {current}
device                    partition=C:
path                      \Windows\system32\winload.exe
description                Windows Server 2012
locale                    en-US
inherit                    {bootloadersettings}
recoverysequence           {424ad145-811e-11e2-abd4-ee945fee3b56}
recoveryenabled            Yes
allowedinmemorysettings    0x150000075
osdevice                   partition=C:
systemroot                 \Windows
resumeobject               {424ad143-811e-11e2-abd4-ee945fee3b56}
nx                         OptOut
hypervisorlaunchtype       Auto

```

Launch Server Manager. Under Tools, choose Hyper-V Manager and navigate to your server. Note that there are no VMs. However, if you click the Virtual Switch Manager action, you'll see a single virtual

**Figure 4**

Changing the Virtual Switch Name and Sharing It with the Management OS

switch that has the name of the network adapter controller; for example, Realtek PCIe GBE Family Controller, as shown in Figure 4.

I recommend renaming the virtual switch something useful, such as External Switch, to represent the network to which it connects. Using consistent naming for switches across your Hyper-V hosts is important: If you move VMs between hosts, a switch of the same name must exist on both the target and source hosts if the VM is to maintain its network connectivity. Also as Figure 4 shows, clear the check box for the *Allow management operating system to share this network adapter* option. That option is needed only if you don't have a separate network adapter for management of the host or if you have only one network adapter that's shared for VM and host traffic. You can also use this interface to create additional switches, as required.

You are now ready to start creating VMs on your standalone host. For maximum capability, cluster multiple Hyper-V hosts and enable live migration. (I go through the complete live migration setup process in

the article “[Shared-Nothing VM Live Migration with Windows Server 2012 Hyper-V.](#)”)

There are likely some other steps that you should at least consider on your new Hyper-V server:

- If you enabled Windows Update, you probably don’t want it to automatically reboot your server. If you’re using an enterprise patch-management solution, make sure to define a maintenance window, outside of business hours, during which your Hyper-V server can reboot. While your Hyper-V server reboots, all your VMs will be unavailable.
- If you run malware protection on your Hyper-V server, you should exclude certain files and folders from scanning, for performance reasons. These exclusions are documented in the Microsoft article “[Virtual machines are missing, or error 0x800704C8, 0x80070037, or 0x800703E3 occurs when you try to start or create a virtual machine.](#)”
- Make sure to back up your VMs. The good news with Windows Server 2012 is that as long as your VMs are backed up, importing those VMs from backup into a new Hyper-V server is easy. You don’t need to export the VMs first.
- If you have a Hyper-V server with a large amount of memory, then Windows by default creates a large pagefile. This file can be manually reduced, because the VMs use most of the memory. You can safely create a 4GB pagefile for the Hyper-V host, using the Control Panel System applet. (Go to the Advanced tab, click the Performance Settings button, go the Advanced tab again, click the Change button under Virtual Memory, then set a custom size. Click the Set and OK buttons on all the open dialog boxes.)

## A Whirlwind Tour

This has been a whirlwind tour of Hyper-V installation. The process is simple, but remember these considerations when choosing your hardware and your configuration level. In a future article, we’ll look at the process of creating VMs on your new host. ■

# Expanding Tabs to Spaces in PowerShell

## A custom function makes this task easy

**T**he tab character has a long history in computing. Tabs were introduced in typewriters, where typists could specify one or more tab stops on the page. Pressing the Tab key would advance the carriage to the next tab stop. In ASCII code on computers, character 9 is designated as the tab. When displaying a tab character in a teletype-like display (e.g., UNIX terminal, Windows console program), the computer will advance the cursor to the next column that's a multiple of eight, where the count starts at column 0. For example, if the cursor is in any column from column 0 through column 7, a tab will advance the cursor to column 8 (which is really the ninth column because the computer is counting from column 0).

Tab characters are also used in other ways in computers. For example, various database and spreadsheet tools let you output data in tab-separated values (TSV) format, where tab characters separate the data items in each row. In addition, scripters and programmers have long debated amongst themselves about whether they should indent code using tabs or spaces. Both techniques have their advantages, but one thing is for sure: You can't tell whether a file contains spaces or tab characters using the `Cmd.exe Type` command, the [Windows PowerShell Get-Content](#) cmdlet, or Notepad because the tabs will appear as spaces.

To prevent confusion, it's often helpful to "de-tab" the contents of a file—that is, expand the tabs to the correct number of spaces. I like to do this for text files in which the tab characters are used for indenting, such as scripts, XML files, and HTML files. Although the `More.com` program in Windows can expand tabs to spaces, I created



### Bill Stewart

is a scripting guru who works for Indian Health Service in Albuquerque, New Mexico. He's a contributing editor for *Windows IT Pro* and a moderator for Microsoft's Scripting Guys forum. He offers free tools on his website.



**Email**



**Website**

a native PowerShell function named `Expand-Tab` to perform this task so that I could take better advantage of PowerShell's pipeline.

## Introducing the Expand-Tab Function

Listing 1 shows the short but handy `Expand-Tab` function. For each line of input it receives, the function uses a regular expression to output the line with the tab characters replaced by the appropriate number of spaces. You can even specify the number of spaces you want to use for each indent (8 by default) or 0 if you want to remove the tab characters altogether. Let's take a look at how this works.

Listing 1: The `Expand-Tab` Function

```
function Expand-Tab {
    param([UInt32] $TabWidth = 8)
    process {
        $line = $_
        while ( $TRUE ) {
            $i = $line.IndexOf([Char] 9)
            if ( $i -eq -1 ) { break }
            if ( $TabWidth -gt 0 ) {
                $pad = " " * ($TabWidth - ($i % $TabWidth))
            } else {
                $pad = ""
            }
            $line = $line -replace "^(^\\t){$i})\\t(.*)$",
                "`$1$pad`$2"
        }
        $line
    }
}
```

The `Expand-Tab` function uses a *process* script block to do something to each line of input it receives. First, the function assigns the variable *\$line* to each input line (i.e., *\$\_*). Then, it uses a *while* loop

that repeats until the input line doesn't contain any tab characters. The `$i` variable contains the position in the string where the tab character occurs. If `$i` is -1 (i.e., no tab character), the function uses the `break` statement to exit from the `while` loop.

Next, the function checks whether `$TabWidth` is greater than 0. If it is, the function creates a string, `$pad`, that contains the needed number of spaces using PowerShell's `*` operator. In PowerShell, `string * n` means "output *string* concatenated *n* times," so `$pad` will contain `$TabWidth - ($i % $TabWidth)` spaces. If `$TabWidth` is 0, `$pad` is set to "" (i.e., an empty string).

Finally, the function uses the `-replace` operator, which uses a regular expression to output a copy of `$line` with the tab characters replaced by `$pad` (i.e., the calculated number of spaces). Table 1 explains the components of the regular expression.

**You can't tell whether a file contains spaces or tab characters using the `Cmd.exe Type` command, the PowerShell `Get-Content` cmdlet, or Notepad because the tabs will appear as spaces.**

Table 1: Regular Expression Components

| Pattern   | Meaning   |
|---|---|
| <code>^</code>  | Find beginning of string  |
| <code>([^\t]{<i>\$i</i>})</code>  | Not a tab character, <i>\$i</i> times; ( ) = first group (i.e., <i>\$1</i> in the replacement string) |
| <code>\t</code>   | A tab character   |
| <code>(.*)</code>   | Any character, 0 or more times; ( ) = second group (i.e., <i>\$2</i> in the replacement string)       |
| <code>\$</code>   | Find end of string  |
| <code>`\$1</code>   | Replace with first group <sup>†</sup>   |
| <code>\$pad</code>  | Replace with calculated number of spaces  |
| <code>`\$2</code>   | Replace with second group <sup>†</sup>  |
| <sup>†</sup> The backtick ( <code>`</code> ) character is needed in the replacement expression to prevent PowerShell from interpreting <i>\$1</i> or <i>\$2</i> as a variable name. |   |

## Using the Expand-Tab Function

I put the `Expand-Tab` function in my PowerShell profile so that it's always available. Here's an example of how to use it:



```
Get-Content t1.ps1 | Expand-Tab | Out-File t2.ps1
```

This command will get the contents of the t1.ps1 file, expand each tab to eight spaces, and save the “de-tabbed” contents in a file named t2.ps1. If the default tab width of eight spaces is too wide, you can specify a different tab width. For example, if you prefer two spaces, you’d use the command:

```
Get-Content t1.txt | Expand-Tab 2 | Out-File t2.txt
```

Note that you don’t need to specify the -TabWidth parameter name. PowerShell knows that the function’s first parameter is -TabWidth.

### Take Control of Your Tabs

By adding the Expand-Tab function to your PowerShell profile, you’ll no longer have to worry about whether your text files contain spaces or tabs. You can download the code for the Expand-Tab Function by clicking the Download button. ■

**Download**



Download the code

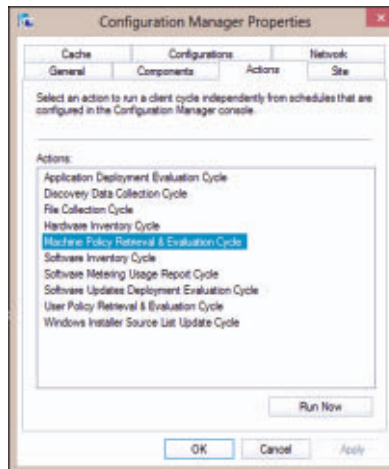
# Understanding System Center 2012 Configuration Manager

## Manage nearly every aspect of client computer configuration

**S**ystem Center Configuration Manager 2012 SP1 is Microsoft's solution for managing the configuration of client computers and devices, from computers running Windows XP through Windows 8 to systems running Mac OS X, Windows Phone, iOS, and Google Android. You can use Configuration Manager to perform tasks including software deployment, software update deployment, anti-malware client management, OS deployment, and hardware and software inventory. In this article, I present some core Configuration Manager concepts to give you a greater understanding of the product's functionality.

### Configuration Manager Client

Configuration Manager performs most operations through special client software installed on each device. You can install the Configuration Manager client directly from the Configuration Manager console, manually using a traditional installer, or baked into an OS image. The client performs tasks such as retrieving policies from Configuration Manager and performing software and hardware inventory, software installation, software updates, and software metering. Figure 1 shows the Configuration Manager client.



### Orin Thomas

is a contributing editor for *Windows IT Pro* and a Windows Security MVP. He has authored or coauthored more than a dozen books for Microsoft Press.



Email



Blog

**Figure 1**  
The Configuration Manager Client

## Operating System Deployment

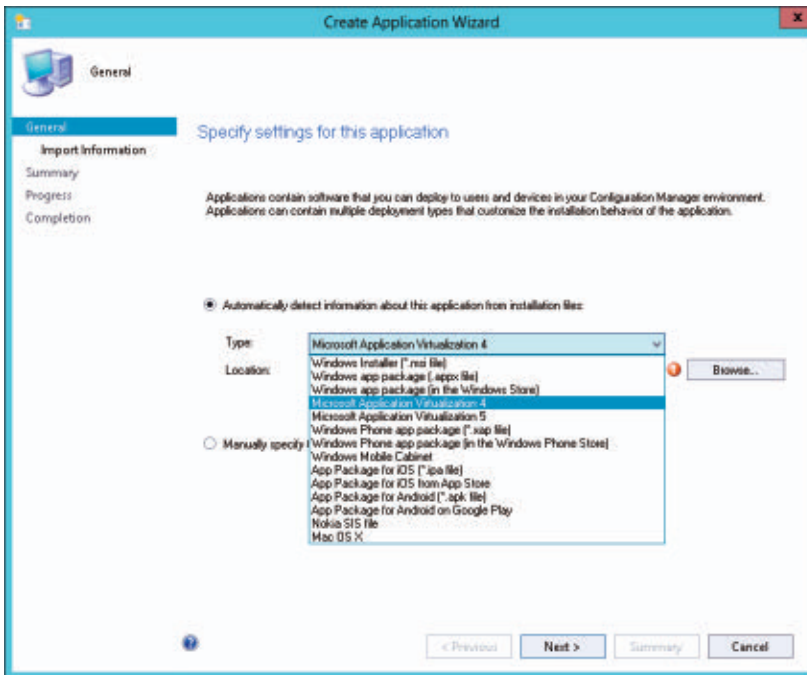
Configuration Manager integrates with Windows Deployment Services to allow you to perform OS deployment and image capture. You automate these processes by configuring task sequences. For example, you can configure task sequences to automate the deployment of a reference image, populate that image with applications and updates, and then perform an image capture of the deployed reference computer. Configuration Manager also supports the automation of offline image maintenance, so you can keep images up to date with the latest patches, hotfixes, and service packs without having to perform full deployment and capture operations.

## Software Deployment

Most organizations use Configuration Manager primarily as a software deployment solution. Configuration Manager 2012 supports software deployment using Packages and Programs, which lets you use packages and programs from Configuration Manager 2007 with Configuration Manager 2012. Configuration Manager applications are a new way of performing software deployment operations that give you substantially more options than the traditional Packages and Programs method. You can use Configuration Manager to deploy software in the formats that Figure 2 shows.

One of the primary benefits of using Configuration Manager applications for software deployment is that they support multiple deployment types. A deployment type lets you deploy an application in a different way, depending on the properties of the computer to which you're deploying the application. For example, you could configure an application so that it's installed as a .msi file on a computer that has one set of properties (e.g., a particular CPU or amount of RAM) and as a .appx file on a computer that has another set of properties.

Configuration Manager 2012 also allows a user's primary device to be set, so you can use that as a condition when configuring a deployment type. For example, you might use this setting to deploy an application



**Figure 2**  
Software Formats

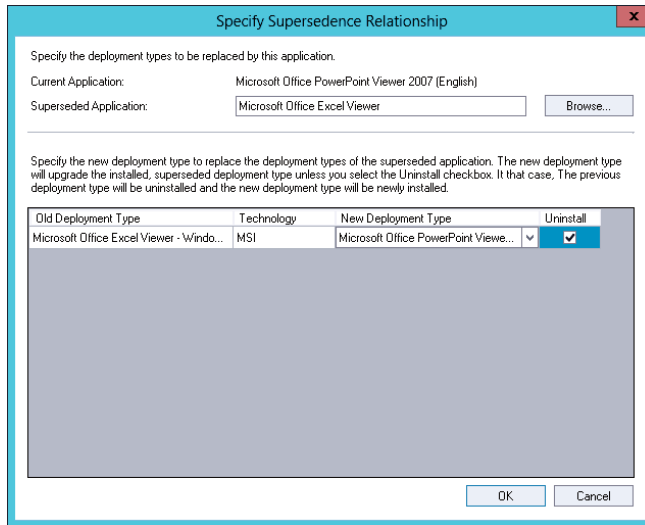
sequenced using Microsoft Application Virtualization (App-V) if you need to deploy an application to a user who isn't signed on to a computer designated as his or her primary device, and deploy the application as a .msi if the user is signed on to a computer designated as the primary device.

Configuration Manager software deployment also supports the following additional features:

- **Application Catalog**—This self-service portal allows users to request software that they can either install automatically or install subject to approval. Deploying this feature lets users request software themselves without having to lodge a Help desk ticket.
- **Application dependencies**—When properly configured, dependencies allow other Configuration Manager applications to be deployed to support the application (e.g., having Configuration Manager automatically deploy the App-V client to a computer before deploying a virtualized application).

- **Application Supersedence**—This feature, which Figure 3 shows, lets you configure existing deployment types so that they're replaced. You can use this to upgrade from one version of an application to another automatically, or to replace one application with another automatically.

**Figure 3**  
Supersedence



## Software Updates

Configuration Manager builds on Windows Server Update Services (WSUS), and WSUS is a necessary component in a software update point. Configuration Manager 2012 provides you with more options than WSUS, allowing the deployment of updates for Microsoft products as well as allowing the deployment and management of software updates for third-party products.

Configuration Manager 2012 provides sophisticated options for the automatic deployment of updates. You can also use Configuration Manager 2012 to perform offline updates to OS deployment images. Configuration Manager 2012 includes a number of sophisticated reports, which provide software update administrators a more accurate picture of how successful software update deployment has been in comparison with WSUS.

## Inventory and Metering

Configuration Manager lets you perform software and hardware inventory, as well as perform software metering. Software and hardware inventory are straightforward. The Configuration Manager client does an audit of all the hardware and software on the device and reports that information back to Configuration Manager. You can use the information generated by this inventory to create Configuration Manager queries, upon which you can base device collections.

Software metering lets you track how often particular applications are used. For example, metering allows you to determine how often a particular CAD program is run. With Configuration Manager, you can enable the automatic creation of metering rules in a disabled state for any application that's used on a specified percentage of computers in the organization. You can then enable these metering rules to track utilization of a specific application.

## Compliance

Compliance lets you set a configuration baseline (e.g., registry setting, existence of a file or folder, or particular version of an application) against which you can measure computers in your organization. You use the compliance functionality of Configuration Manager to determine whether the configuration of those computers meets or falls below your organizational needs. You can use compliance to ensure that computers in your organization meet legislative or security requirements.

## Endpoint Protection

System Center Endpoint Protection is an evolution of Microsoft's Forefront Endpoint Protection anti-malware product. Rather than having the anti-malware aspect of client health managed through a computer running the Forefront console and other parts of client health (e.g., adherence to a configuration baseline) monitored in a separate console, the integration of Endpoint Protection lets you monitor all aspects of client health through Configuration Manager.

As the information from the Endpoint Protection client is stored in the Configuration Manager database, it can also be used as the foundation for query-based Configuration Manager device collections.

## **Windows Intune Integration**

With the release of System Center 2012 SP1, you can integrate an on-premises Configuration Manager deployment with a Windows Intune subscription. Rather than set up Configuration Manager on the perimeter network so that clients on the Internet are able to contact the servers directly, integrating Windows Intune with Configuration Manager allows Windows Intune to take on the role of managing devices outside the perimeter network. When you integrate Configuration Manager with Windows Intune, you can use the Configuration Manager console to manage non-domain-joined devices, including mobile devices, that are connected to networks both inside and outside the perimeter network.

## **User and Device Collections**

You perform actions in Configuration Manager, such as software deployment, against collections. Configuration Manager supports user collections, made up of user accounts and groups, and device collections, made up of computers and other devices. You can create collections manually, or you can configure collections to be generated dynamically—for example, a device collection that includes only computers that have Microsoft Office 2010 SP1 installed, or users who report to a specific manager.

## **Manage All Aspects**

Configuration Manager 2012 SP1 lets you manage almost all aspects of client computer configuration. Not only does it provide a platform for the deployment and management of applications, but you can also use it as an anti-malware, software update, and configuration monitoring solution. ■



# FAQ

## Answers to Your Questions

**Q:** How do I fix an error in Server Manager for Windows Server 2012 performance counters?

**A:** Windows Server 2012 boxes have a user-defined data collector set with the performance counters Server Manager needs, but it's typically not started. To start it, and to gather performance information, select *Start Performance Counters in Server Manager*. If you receive an error in Server Manager—"Performance counter refresh failed with the following error: Data Collector Set was not found"—the user-defined data collector set is missing.

Log on to the server causing the error. Launch Performance Monitor (perfmon.exe) and under Data Collector Sets, User Defined, look for Server Manager Performance Monitor. Right-click it and select Start. If it's not present, log on to a server whose performance counters can be monitored. Open an elevated command prompt and export the Server Manager Performance Monitor data collector set:

```
logman export -n "Server Manager Performance Monitor" -xml
SMPM.xml
```

Open the generated XML file and remove the section `< Security > List of IDs </Security >`, then save the file. Open an elevated command prompt where the data collector set is missing and import the XML file, then start the performance counters:

```
logman import -n "Server Manager Performance Monitor" -xml
SMPM.xml
```



**John Savill**



**Jan De Clercq**

logman start "Server Manager Performance Monitor"

Refresh the server in Server Manager. Performance data should be available now.

—John Savill

**Q:** Can multiple NFS or SMB shares point to the same base folder?

**A:** Yes. There's no restriction on the folders pointed to by the SMB or NFS share. Both SMB and NFS are file-based sharing protocols and allow access to the data stored within the folder and subfolders the share points to.

These protocols don't change the underlying file system. Thus, having multiple shares, even different protocol shares, that point to the same base folder isn't a problem.

—John Savill

**Q:** How do I force a discovery of a System Center 2012 Data Protection Manager server?

**A:** The easiest way to force a discovery in System Center 2012 Operations Manager of Data Protection Manager (DPM) is to restart the health service (System Center Management) on the DPM server:

```
net stop healthservice
net start healthservice
```

Then look at the Operations Manager event log for Event ID 125:

```
Get-EventLog -LogName 'Operations Manager' -Newest 6
-InstanceId 125 | ft TimeGenerated, Message
```

This category of IDs shows the detail of the discovery start, process, and completion (see Figure 1 for an example of output).

| TimeGenerated         | Message  |
|-----------------------|--|
| -----                 | -----  |
| 3/22/2013 2:30:25 PM  | DPMServerDiscovery : DPM server discovery completed... |
| 3/22/2013 2:30:21 PM  | DPMServerDiscovery : DPM server discovery SCOM Mana... |
| 3/22/2013 2:30:21 PM  | DPMServerDiscovery : DPM server discovery start event  |
| 3/22/2013 11:49:05 AM | DPMServerDiscovery : DPM server discovery completed... |
| 3/22/2013 11:49:01 AM | DPMServerDiscovery : DPM server discovery SCOM Mana... |
| 3/22/2013 11:49:01 AM | DPMServerDiscovery : DPM server discovery start event  |

**Figure 1**

Operations Manager  
Event Log for Event  
ID 125

—John Savill

**Q:** Can you provide a short list of the most important tools I can use to troubleshoot a Network Access Protection problem?

**A:** For Network Access Protection (NAP) troubleshooting on the server side, you should first check the NAP-specific error messages found in this Event Viewer container: Custom Views\Server Roles\Network Policy and Access Services. To view NAP configuration information on a NAP server, you can use the following netsh commands:

- For NAP Network Policy Server (NPS) configuration information:

```
netsh nps show config
```

- For NAP Health Registration Authority (HRA) configuration information:

```
netsh nap hra show config
```

For NAP troubleshooting on the client side, check for error messages in the following Event Viewer container: Applications and Services

Logs\Microsoft\Windows\Network Access Protection\Operational. To view NAP configuration information on a client, you can use the following netsh commands:

- For client Group Policy configuration:

```
netsh nap client show group
```

- For client local policy configuration:

```
netsh nap client show config
```

- For client NAP state:

```
netsh nap client show state
```

To see which NAP System Health Agent (SHA) is causing problems, use the NAP-related events in the Event Viewer. These events contain an error message with an identifier of the SHA that caused the error. The HKEY\_LOCAL\_MACHINE\SYSTEM\CurrentControlSet\Services\napagent\Shas registry container holds a list of all SHAs that are active on your system.

For more about NAP-specific events and their IDs, see “NAP event logs” in the Microsoft article “[Tools for Troubleshooting NAP](#).” For more about the event IDs related to NAP agent communication with the SHA, see “[NAP Agent Communication with the SHA](#).”

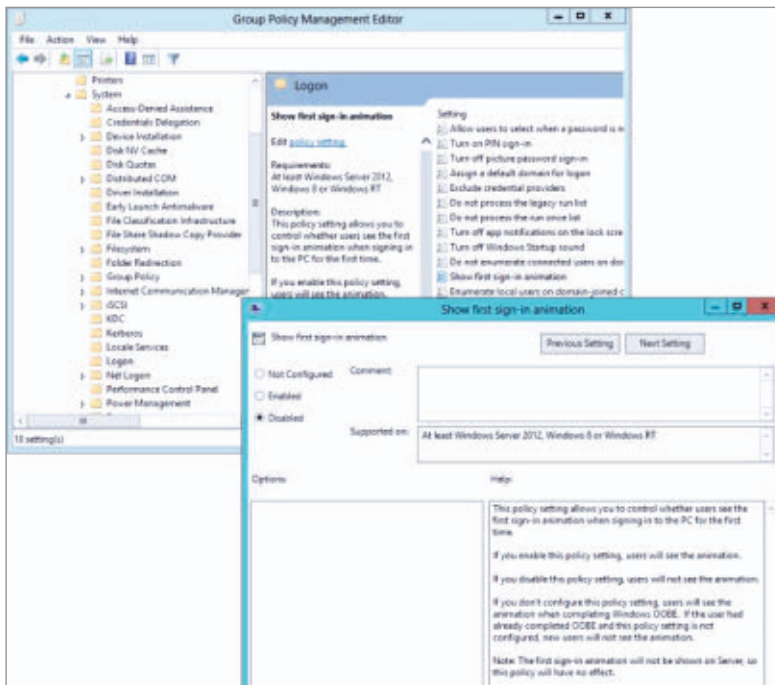
If you have Microsoft System Center Configuration Manager (SCCM) installed, you can use it for advanced logging and data collection on NAP clients. To learn more about SCCM NAP-specific log files, see “[Log Files for Network Access Protection](#).” For more about NAP, see my article “[A Microsoft Network Access Protection \(NAP\) Primer](#)” and Russell Smith’s “[Managing Security Dependencies on Windows Networks](#).”

—Jan De Clercq

## Q: How can I disable the Windows 8 starting animation shown for new users?

**A:** Typically, the first time a user logs on to Windows 8, he or she is presented with the Windows 8 sign-in animation. It shows the basic new gestures for Windows 8. You can disable it by using Group Policy:

1. Open an existing Group Policy Object (GPO) in Group Policy Management Editor (see Figure 2) or create a new one.



**Figure 2**  
Changing Windows 8  
Sign-In Animation

2. Navigate to Computer Configuration, Policies, Administrative Templates, System, Logon.
3. Double-click *Show first sign-in animation*.
4. Set the policy to Disabled and click OK, then close the GPO. ■

—John Savill

# Product News for IT Pros



## Spiceworks Partners with Dropbox

Spiceworks announced a technology integration and partnership with Dropbox to help companies simplify the use of Dropbox accounts within their organizations. As part of the partnership, Dropbox features have been integrated into Spiceworks' free IT management application to help streamline how companies discover, manage, and grow Dropbox environments. The Dropbox integration with Spiceworks is designed to help companies and their IT departments quickly understand how often Dropbox is being used within their organizations, view the status of company-issued Dropbox accounts, and monitor allocated space. Spiceworks users will also be able to share and download Dropbox documents from within Spiceworks and assign Help desk tickets to individual Dropbox accounts when employees need help. Current and prospective Dropbox customers can manage their services for free by downloading or updating their Spiceworks applications to the latest version, Spiceworks 6.2. For more information, visit the [Spiceworks website](#).



## NovaStor Protects Windows Server 2012 and Exchange Server 2013

NovaStor launched NovaBACKUP 14.5 for PCs, laptops, and servers. Version 14.5 fully supports current Microsoft technology and offers full data availability for servers running [Windows Server 2012](#) and [Exchange Server 2013](#). With the new Windows 8/UEFI Secure Boot support, NovaBACKUP 14.5 also delivers a comprehensive solution for data backup and restoration, even for future generations of hardware. NovaBACKUP provides a complete product family of scalable data protection for small-to-midsized businesses (SMBs). From single

laptops and servers to Hyper-V or VMware virtual machines (VMs) to entire networks, NovaStor's software protects all your important business data. Wizards guide you through the process of data backup and recovery. NovaBACKUP provides differential and incremental backup, data compression, virus scans, and more. For more information, check out the [NovaStor website](#).

## Dell Embraces Touch

Touch is becoming an increasing reality in the enterprise, so Dell is extending its Precision, Latitude, and OptiPlex brands to the [Windows 8](#) OS with the release of Dell Client Integration Pack 3.1 for Microsoft System Center Configuration Manager (SCCM) 2012 SP1. This new plug-in helps customers seamlessly manage touch devices by enabling remote deployment of Windows 8. Customers no longer have to “hand-install” the OS, but rather easily and automatically deploy a Configuration Manager 2012 console for Dell systems. Version 3.1 of the plug-in will save time and decrease costs with its automation capabilities; let users remotely manage offline systems one-to-many with Dell's Intel vPro extensions for remote BIOS management, battery management, and remote hard drive wipe; and better manage the touch devices that are making their way into the enterprise. For more information, visit the [Dell TechCenter](#).



## A10 Networks Unleashes Thunder

A10 Networks announced its A10 Thunder Series platforms, which are hardware and software Application Delivery Controllers (ADCs) that provide Unified Application Service Gateway functionality by consolidating premium solution modules for intelligent [cloud services](#) in the most efficient form factors. Built on A10's scalable Advanced Core Operating System (ACOS) architecture, the new Thunder Series models consolidate standalone solutions encompassing ADCs and Server Load Balancers, as well as features such as Carrier Grade NAT (CGNAT), IPv6 Migration, DNS Application Firewall, Web Application





Firewall (WAF), SSL Intercept, Distributed Denial of Service (DDoS) protection, Application Access Management (AAM), and more, without license fees. For more information, visit the [A10 Networks website](#).



### **Paessler Announces New PRTG Network Monitor Version**

Paessler announced the availability of its latest update to PRTG Network Monitor. The new version introduces a new Single Page Application (SPA) architecture based on [HTML5](#) for a faster web experience, allowing network administrators to view and digest more information on a single page. Through the use of HTML5, PRTG Network Monitor offers an improved UI. Data presentation is cleaner and easier to understand, with more overall functionality and fewer clicks needed. The improved design allows both the sensor and device pages to offer more information at a single glance. To further improve workflow for network administrators, tasks now take place in pop-up layers within the browser instead of loading new pages, enabling network administrators to easily focus on the task at hand. Users can change the priority and favorite status of objects with just one click. For more information, visit the [Paessler website](#).



### **STORServer Enters the Cloud**

STORServer announced that it's entering the [cloud](#) market with three cloud options for scalable data protection. The cloud solutions are backed by the company's purpose-built data protection appliance and are sold with a data-recovery guarantee, which ensures that customers will not only get their data back in the event of a loss, but the data will also be useable once it is recovered. STORServer offers both private and public cloud offerings for disaster recovery, as well as managed cloud services for automated backup, archive, and disaster recovery. The first option, disaster recovery to a private cloud, is available to existing IBM Tivoli Storage Manager (TSM) and STORServer customers, and it includes a STORServer Backup Appliance for automated disaster recovery that is configured to the enterprise's daily

backup requirements. The second option, disaster recovery to a public cloud, provides a unit-based monthly or annual service charge for copying backup or archive data to a remote location. And the third option—backup, archive, and disaster recovery to a public cloud—is available to any facility and doesn't require a TSM or STORServer installation. Purchase and installation of software clients and database agents, however, is required. For more information, visit the [STORServer website](#).

### **Piriform CCleaner Professional Now Available Through HP**

Piriform announced an agreement with HP to offer its CCleaner Professional computer optimization tool to HP customers. CCleaner Professional will now be available for purchase as a software option with any system sold through HP's North American online Home & Home Office Store. The agreement stands to extend the CCleaner brand to thousands of HP customers. "We're keen to promote CCleaner Professional to more users, and HP is the perfect partner," said Guy Saner, CEO at Piriform. "We're already seeing a very positive response from HP customers and are helping to keep their PCs running at maximum performance." CCleaner helps the user's computer run like new by removing unnecessary files, including unused and temporary files. It clears Internet and download history, eliminating digital "traces" that can compromise privacy and result in [identity theft](#). For more information, check out the [Piriform website](#). ■

Piriform 

# Cutting the Cord

## Product of the Month



**Jason  
Bovberg**

Email



Twitter




Website



Check out what Rob Honeycutt, founder of Timbuk2, has been up to. With the help of a Kickstarter campaign, he is developing the Elroy Bluetooth Earbuds—wearable Bluetooth technology that clips to your clothing, letting you avoid the problems of being tethered to long headphone cords. The earbuds connect to the Bluetooth device with a much shorter cable, and you can connect the buds to the device's magnetic docking station when they're not in use. The successful Kickstarter campaign means the assembly processes are now up and running, so expect to see the Elroy enter the market soon. We look forward to eliminating those tangled cords and feeling even more mobile in this increasingly digital age. Find out more at the [Elroy website](#).



**Figure 1:** Justifiable homicide?



**Figure 2:** Still waiting...

Send us your funny screenshots, oddball product news, and hilarious end-user stories. If we use your submission, you'll receive a *Windows IT Pro* Rubik's Cube.



**Submit**

Search our network of sites dedicated to hands-on technical information for IT professionals.  
[www.windowsitpro.com](http://www.windowsitpro.com)

## Support

Join our discussion forums. Post your questions and get advice from authors, vendors, and other IT professionals.  
[forums.windowsitpro.com](http://forums.windowsitpro.com)

## News

Check out the current news and information about Microsoft Windows technologies.  
[www.winsupersite.com](http://www.winsupersite.com)

## EMAIL NEWSLETTERS

Get free news, commentary, and tips delivered automatically to your desktop.

- Cloud & Virtualization UPDATE
- Dev Pro UPDATE
- Exchange & Outlook UPDATE
- Security UPDATE
- SharePoint Pro UPDATE
- SQL Server Pro UPDATE
- Windows IT Pro UPDATE
- WinInfo Daily UPDATE

## RELATED PRODUCTS

### Windows IT Pro VIP

Get exclusive access to over 40,000 articles and solutions on CD and via the web. Includes FREE access to eBooks and archived eLearning events plus a subscription to either *Windows IT Pro* or *SQL Server Pro*.  
[windowsitpro.com/vip-premium-membership](http://windowsitpro.com/vip-premium-membership)

### SQL Server Pro

Explore the hottest new features of SQL Server, and discover practical tips and tools.  
[www.sqlmag.com](http://www.sqlmag.com)

### Dev Pro

Discover up-to-the-minute expert insights, information on development for IT optimization, and solutions-focused articles at [DevProConnections.com](http://DevProConnections.com), where IT pros creatively and proactively drive business value through technology.  
[www.devproconnections.com](http://www.devproconnections.com)

### SharePoint Pro

Dive into Microsoft SharePoint content offered in specialized articles, member forums, expert tips, and web seminars mentored by a community of peers and professionals.  
[www.sharepointpromag.com](http://www.sharepointpromag.com)

## Advertiser Directory

|   |        |
|---|--------|
| <b>1&amp;1 Internet</b> .....             | 1      |
| <b>Parker Software</b> .....              | 20, 21 |
| <b>PASS Summit</b> .....                  | 12     |
| <b>Western Governors University</b> ..... | 2      |
| <b>Windows IT Pro</b> .....               | 16, 45 |

## Vendor Directory

|                           |        |
|---------------------------|--------|
| <b>A10 Networks</b> ..... | 75, 76 |
| <b>AMD</b> .....          | 47     |
| <b>Android</b> .....      | 6, 11  |
| <b>Apple</b> .....        | 6, 11  |

|                         |                |
|-------------------------|----------------|
| <b>Dell</b> .....       | 75             |
| <b>Elroy</b> .....      | 78             |
| <b>Intel</b> .....      | 10, 43, 45, 47 |
| <b>NovaStor</b> .....   | 74, 75         |
| <b>Paessler</b> .....   | 76             |
| <b>Piriform</b> .....   | 77             |
| <b>Samsung</b> .....    | 11             |
| <b>Spiceworks</b> ..... | 74             |
| <b>STORServer</b> ..... | 76, 77         |